Aquatic Robotics at the University of Southern California

Marine Robotics Research Summer School 2016

Stephanie Kemna <kemna@usc.edu>

July 8, 2016

USC Viterbi School of Engineering

ROBOTIC EMBEDDED SYSTEMS LABORATORY

Including the works of; Arvind Pereira, Jnaneshwar Das, Geoff Hollinger, Andreas Breitenmoser, Artem Molchanov, Lantao Liu, Kai-Chieh Ma, Hordur Heidarsson

Robotic Embedded Systems Lab

Prof. Gaurav Sukhatme

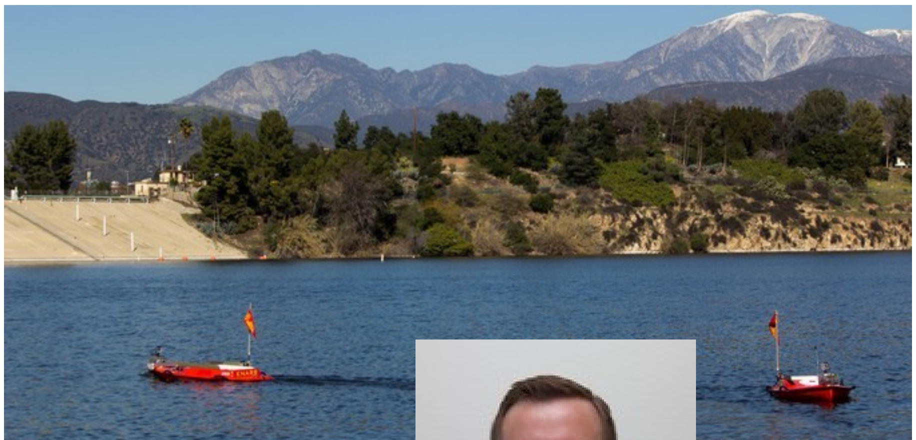

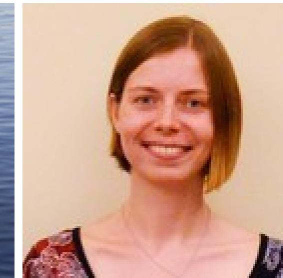

Stephanie Kemna

Andreas Breiten-moser

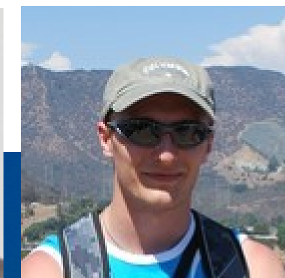

Artem Molchanov

Hörður Heiðarsson

Arvind Pereira

Jnaneshwar Das

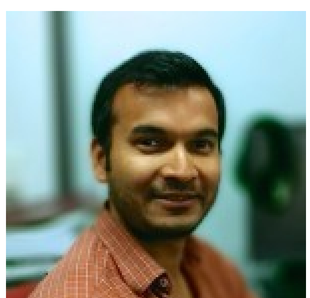

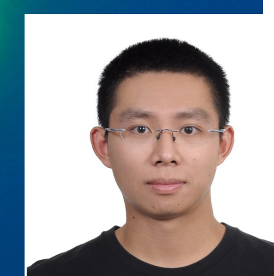

Kai-Chieh Ma

Geoff Hollinger

Lantao Liu

USC Viterbi School of Engineering

ROBOTIC EMBEDDED SYSTEMS LABORATORY

Aquatic Robotics at RESL

- Path planning & adaptive sampling approaches for
  - underwater gliders
  - active drifters
  - autonomous underwater vehicles (AUVs)
- Multi-robot coordination for autonomous underwater and autonomous surface vehicles (ASVs)
- Obstacle avoidance & sensor calibration for ASVs

Path planning for underwater gliders

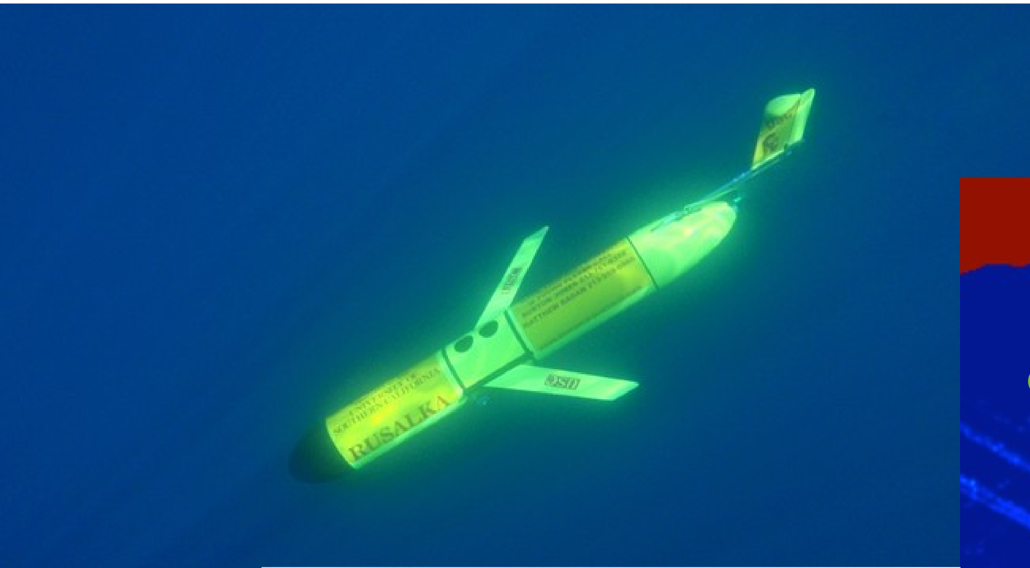

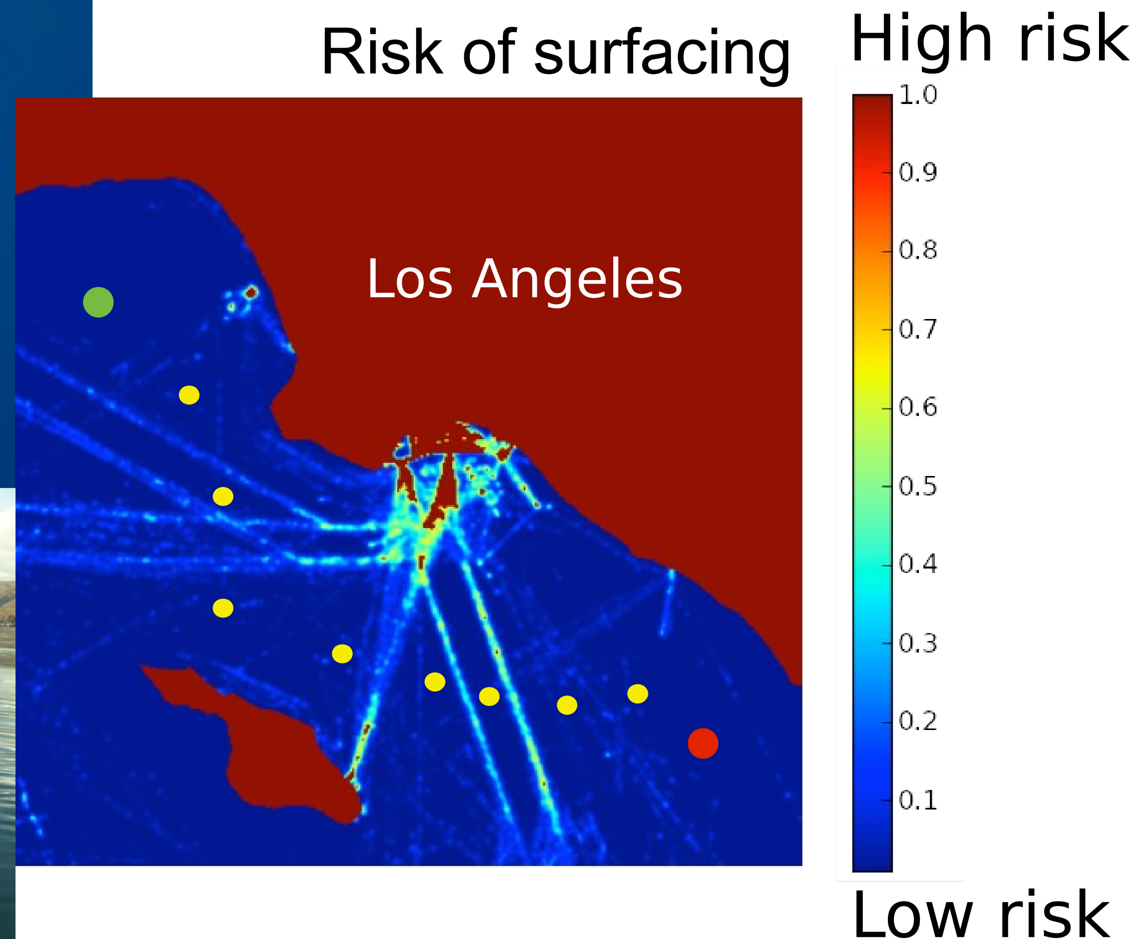

Arvind Pereira, Geoff Hollinger

Slocum gliders

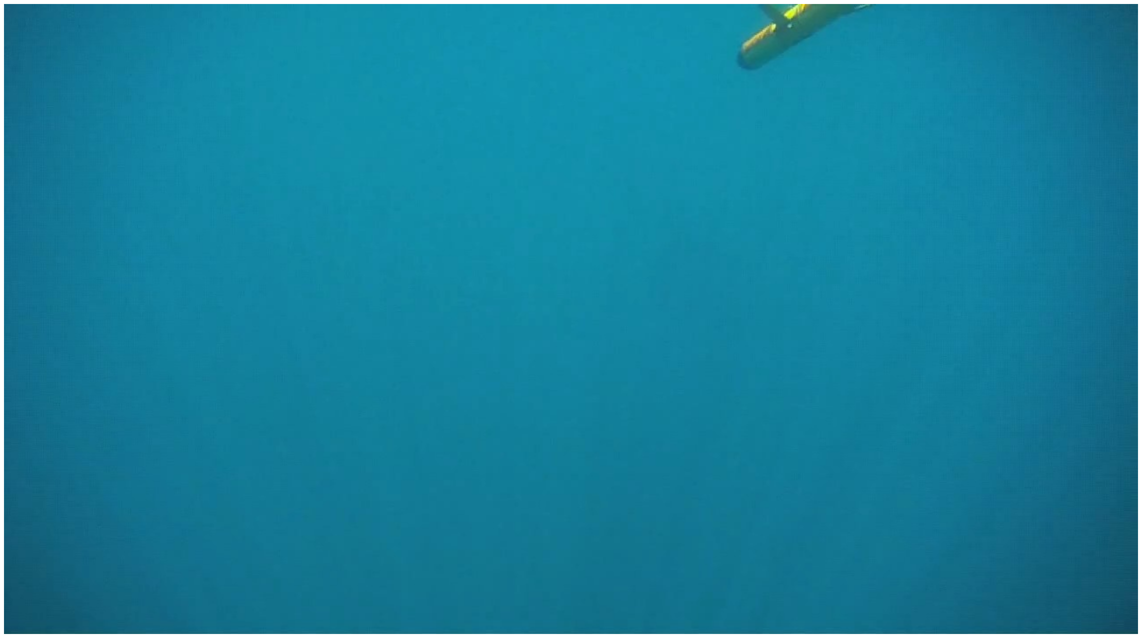

No perception:
No current sensing

Slow moving:
0.3 m/s

Long endurance:
3-4 weeks

Slocum gliders – typical trajectories

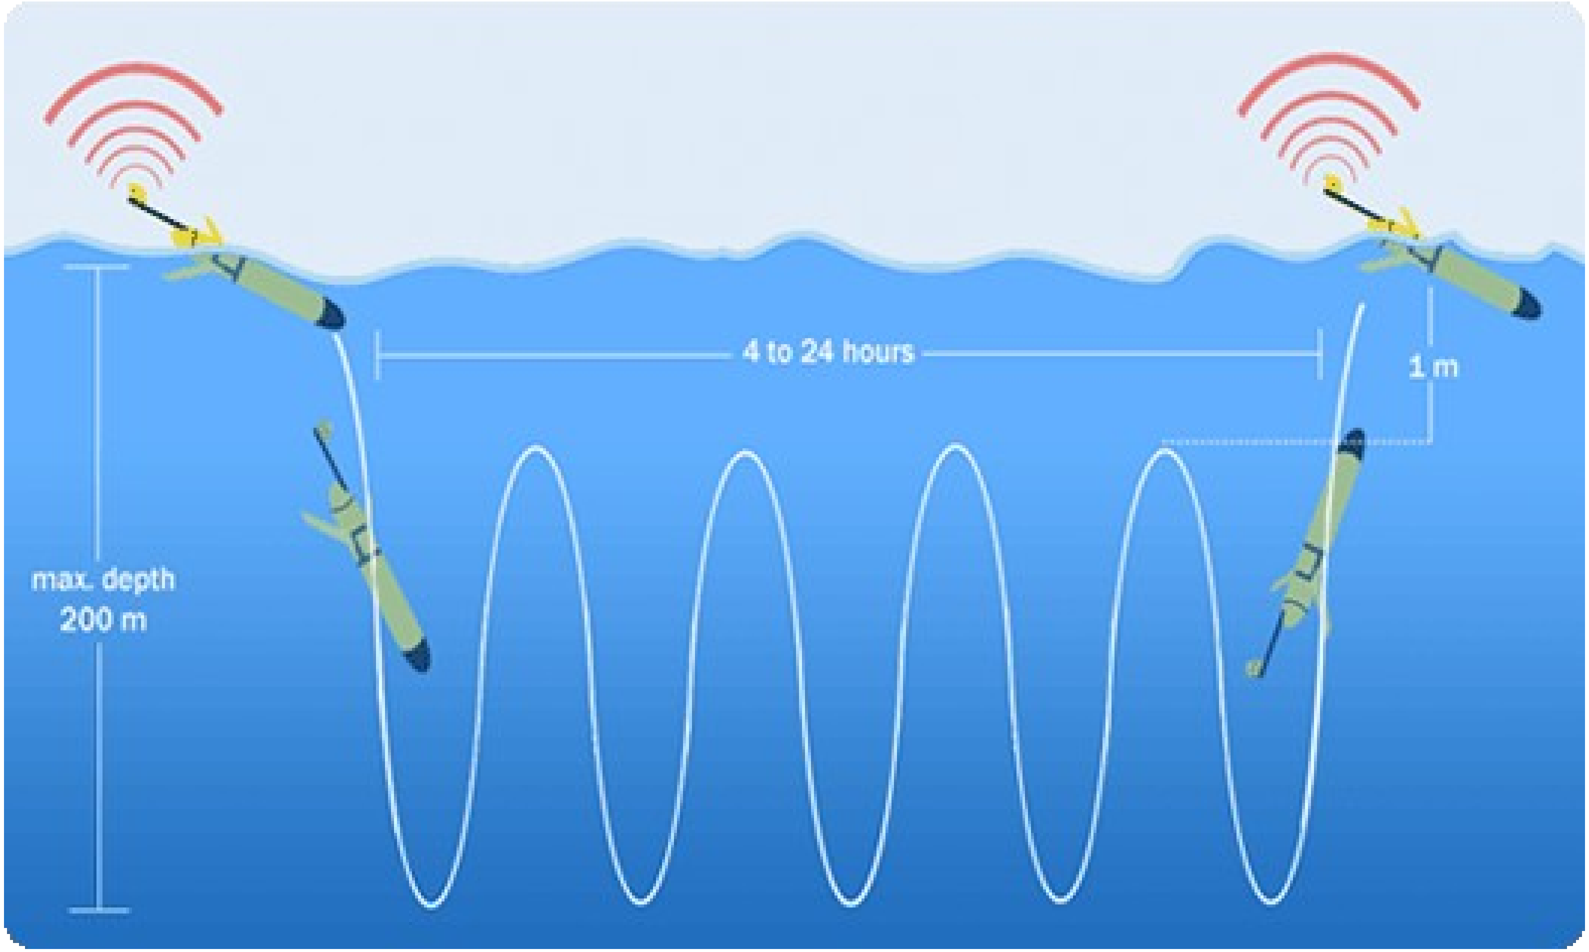

http://www.marine-knowledge.com/wp-content/uploads/2013/10/gliderdiagram.gif

Risk-aware path planning – avoid collisions!

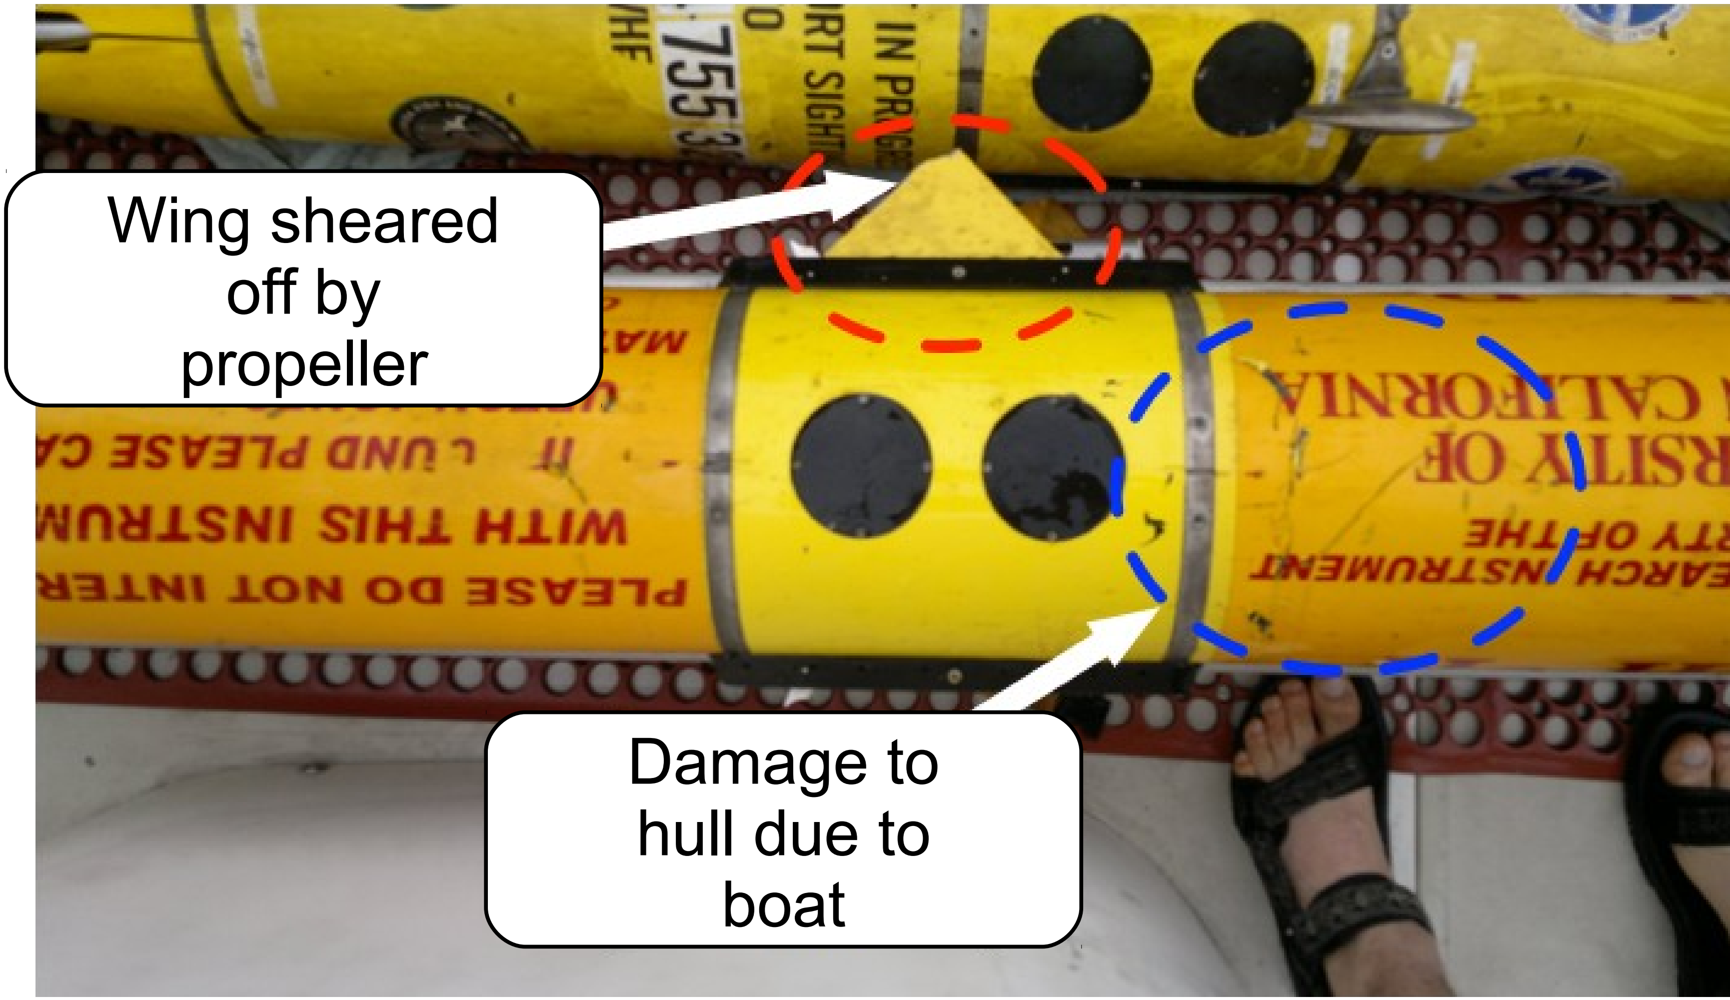

Arvind Pereira, Geoff Hollinger

Picture courtesy of Carl Oberg

Risk-Aware Planning

The probability of collision between ships and AUVs is proportional to ship density [Merckelbach, 2012]

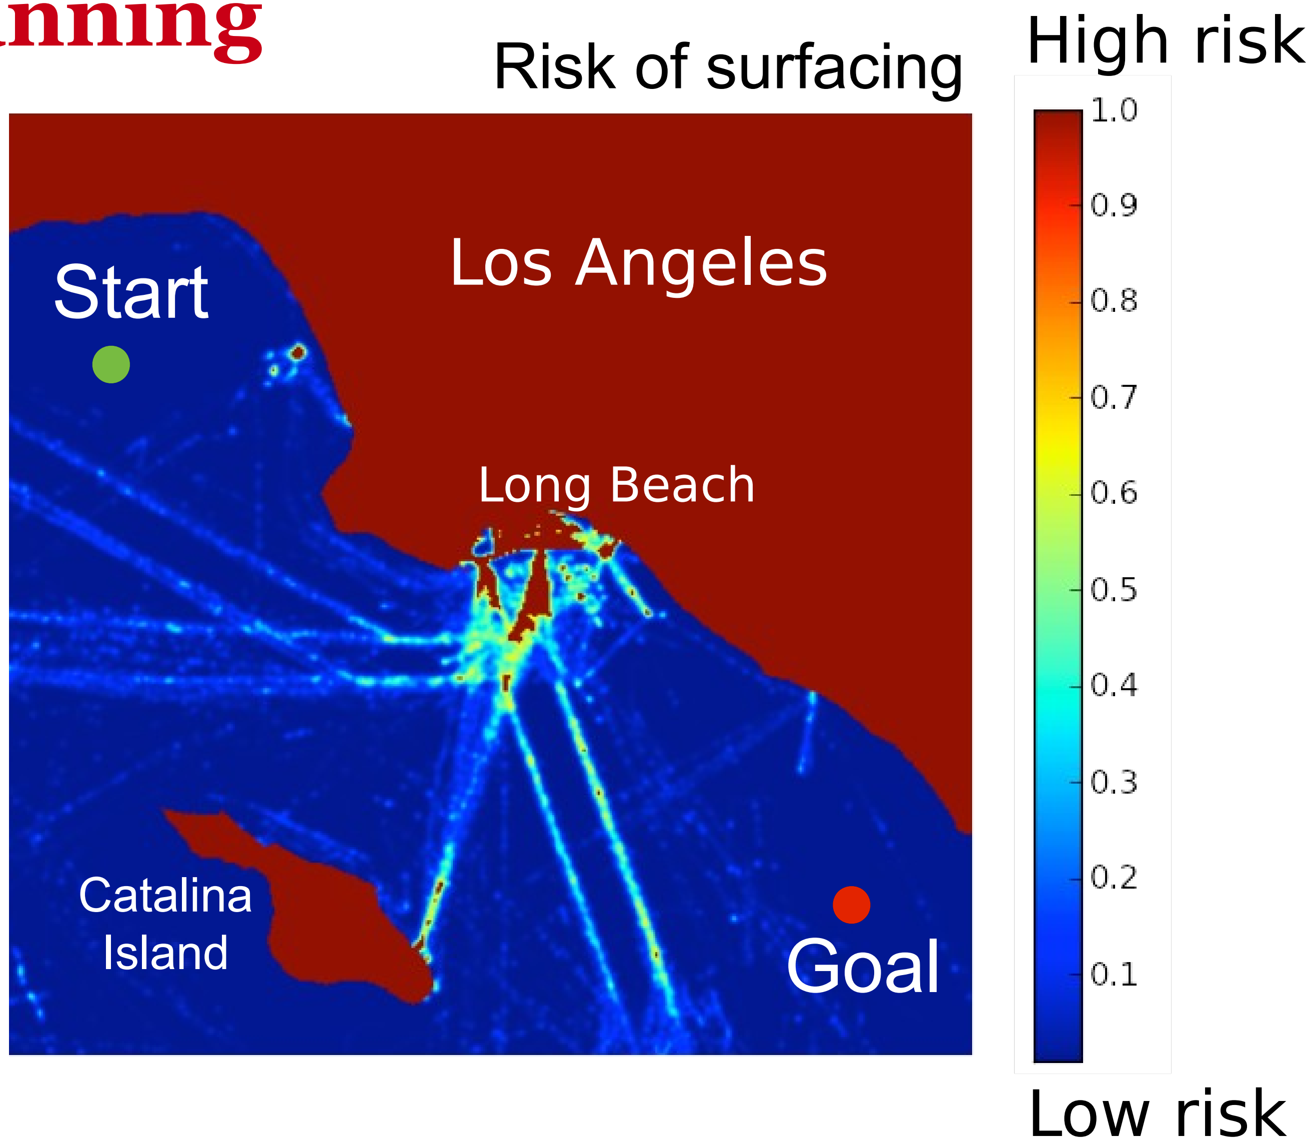

Minimum risk planner

Find path P* with surfacing waypoints *w*:

$$P^* = \underset{P}{\operatorname{argmin}} \sum_i risk(w_i)$$

Subject to constraint:

$$\|e(w_i, w_{i-1})\| \leq d_{max}$$

i.e. max distance between waypoints is limited

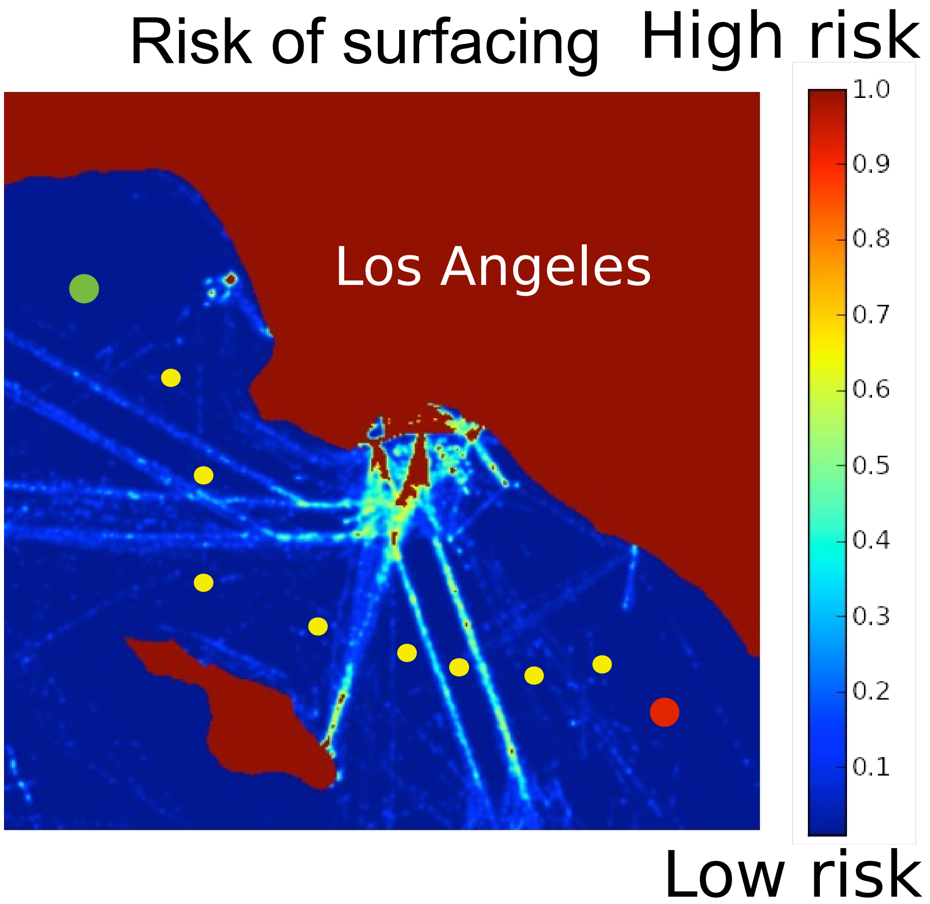

But what if the glider is pushed off course by ocean currents?

But what if the glider is pushed off course by ocean currents?

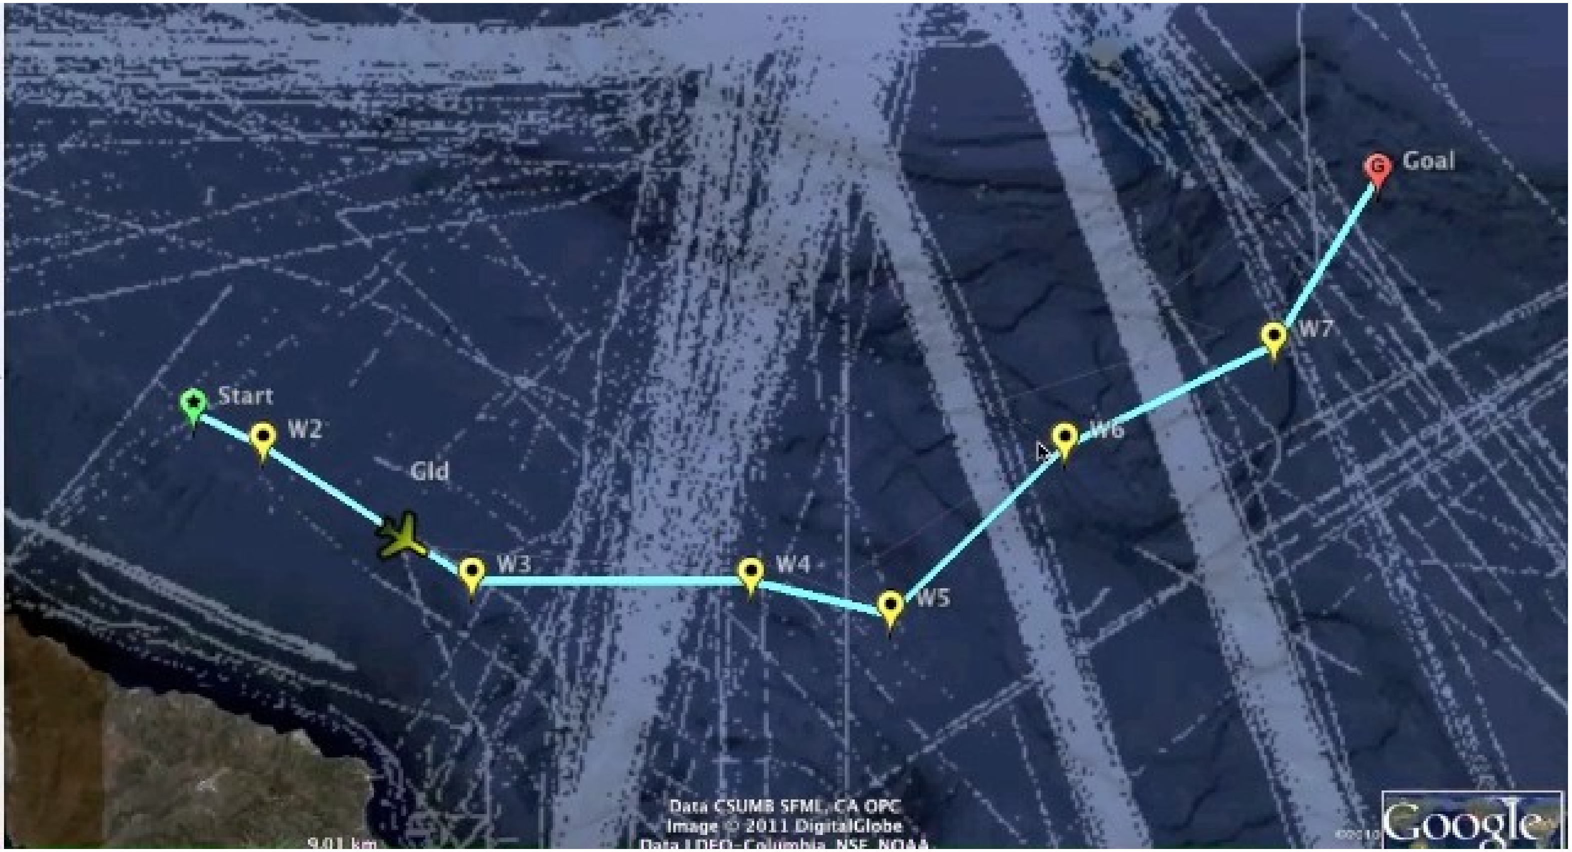

 desired (standard) waypoints

Ocean currents

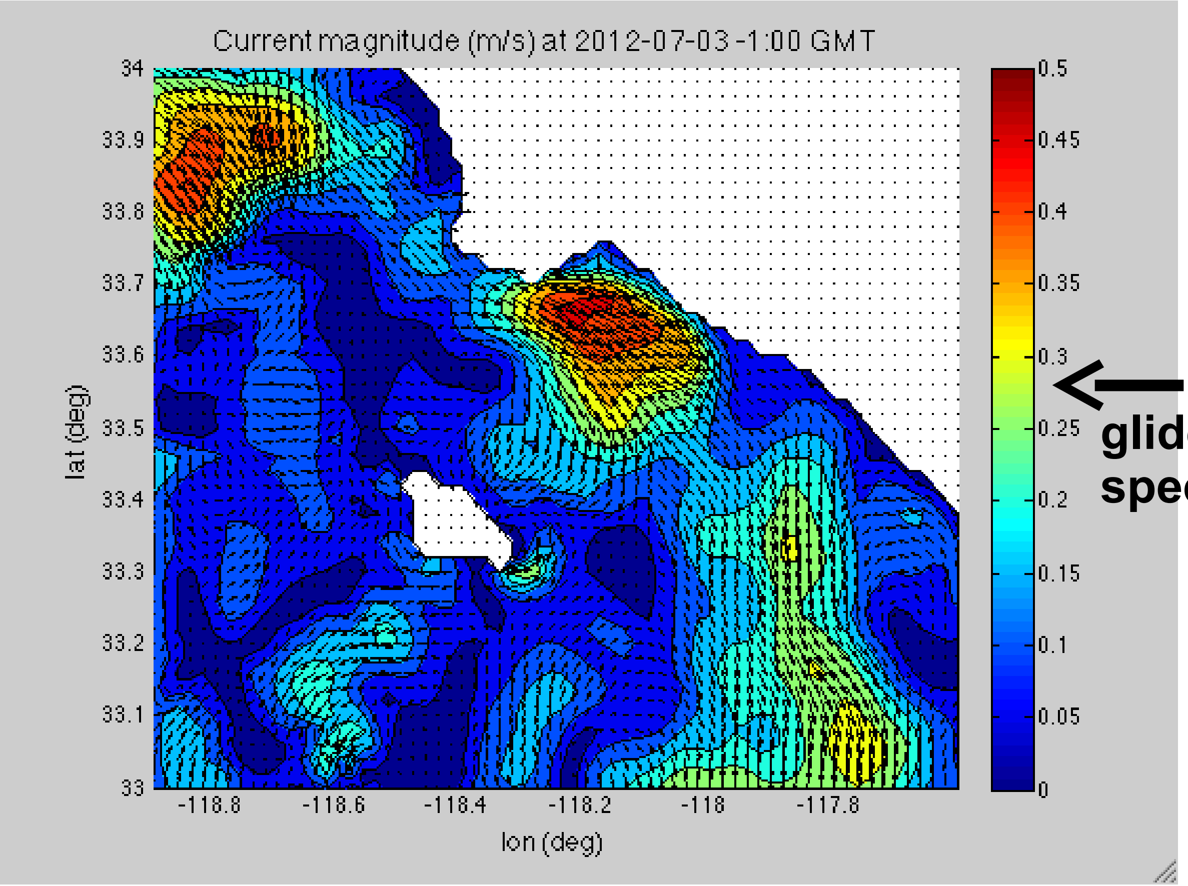

glider speed

- Oceans can have strong currents
- Nearly twice the speed of the glider in red regions
- Direction may change periodically

Incorporating ocean models

Data sources

- HF-radar (surface currents)
- Tide gauges (sea surface height)
- Satellite data (sea surface temperature)
- AUV data
- Etc.

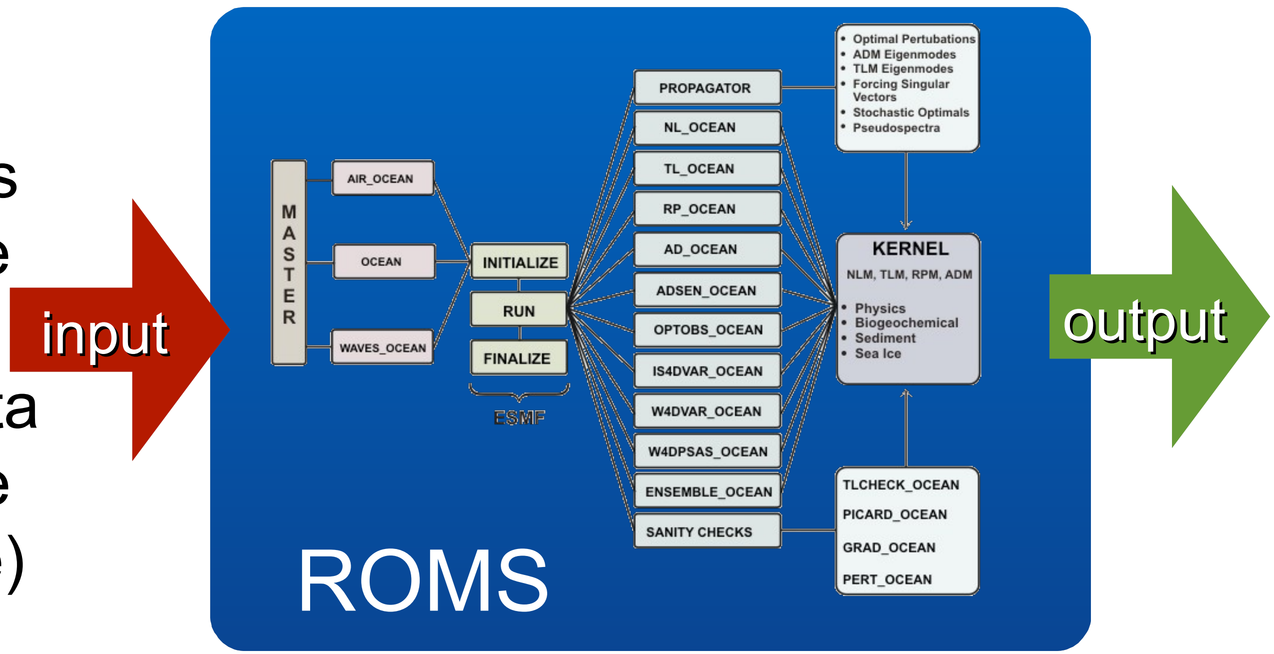

72 hr forecast

- **u - easting**
- **v - northing**
- **w - vertical**
- sal - salinity
- temperature
- sea-surface height

Minimum-Risk planner + pseudo waypoints

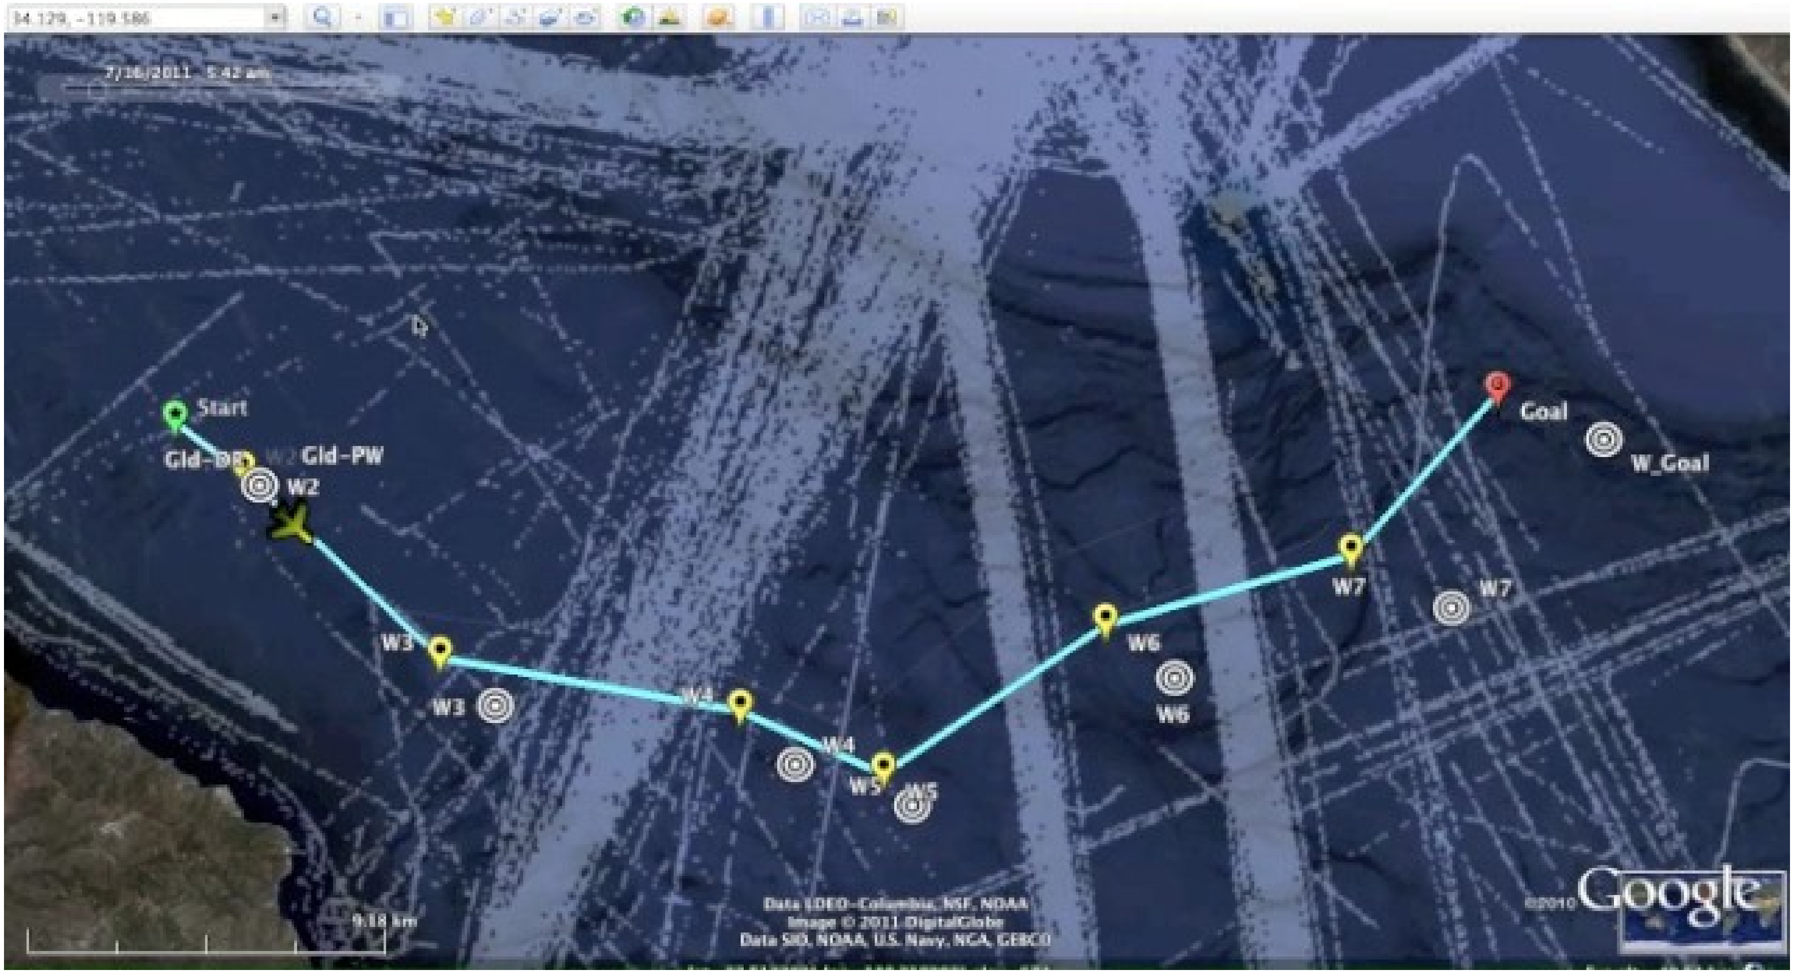

desired (standard) waypoints   pseudo waypoints

But what if the predictions are incorrect?

Ocean current predictions are noisy!

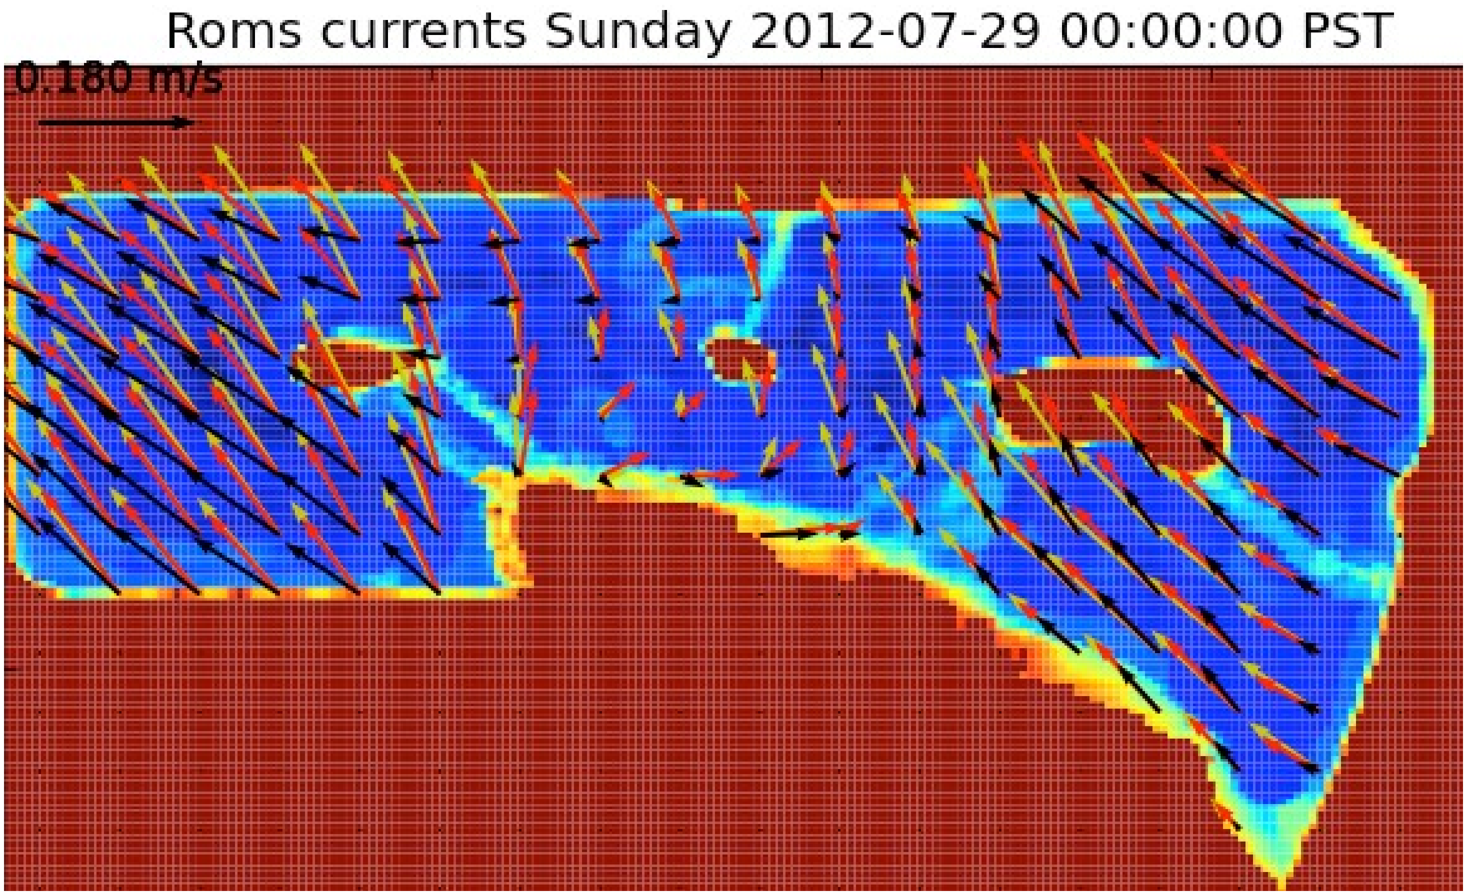

Predicted 48 hrs earlier | Predicted 24 hrs earlier | Nowcast (assimilated)

Oceans currents & prediction uncertainties

Negligible currents (ignore predictions)	Predictable currents
Uncertain predictions (stationary models)	Uncertain predictions (non stationary models)

Path planning for different current systems

Regime	Planner
Negligible currents	Minimum-Risk
Predictable currents	Minimum-Risk planner with pseudo-waypoints
Uncertain (stationary) currents	Minimum Expected Risk planner and risk-aware Markov Decision Process (MDP)
Uncertain (non-stationary) currents	Risk-aware Non-Stationary Markov Decision Process (NSMDP)

Learning better estimates for uncertainty in ocean current predictions

Gaussian Processes: estimate the value **with an uncertainty estimate!**

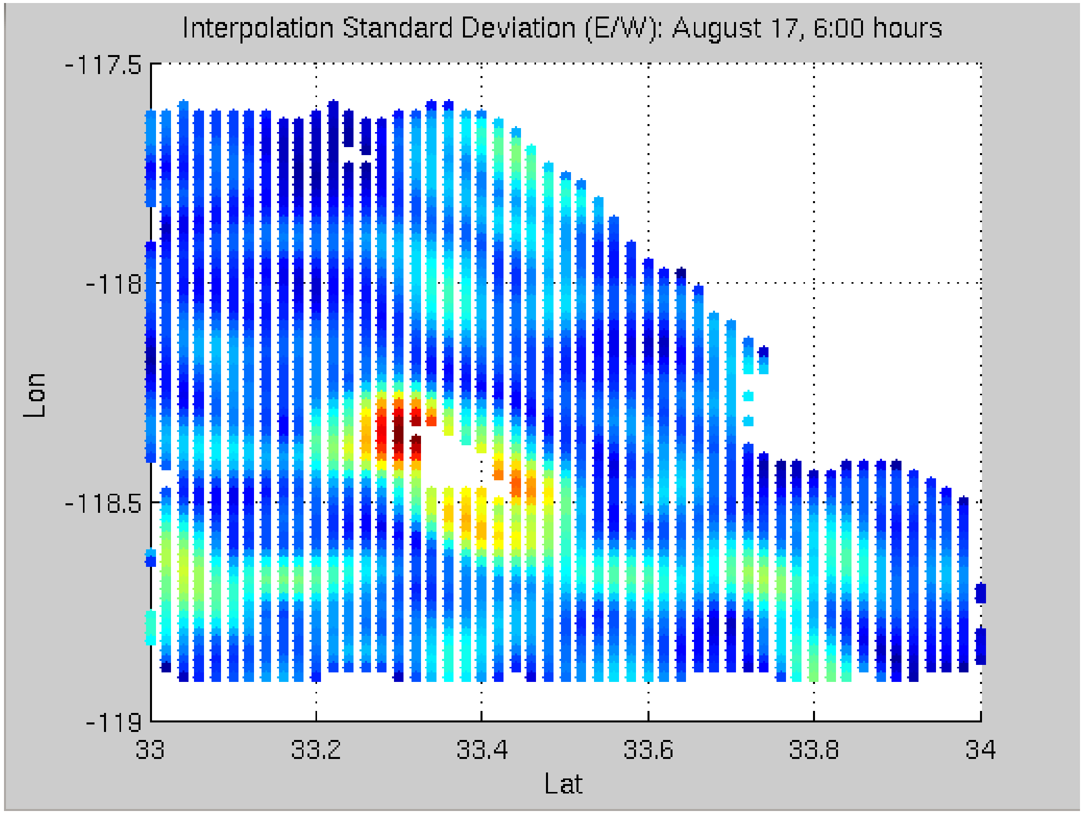

Learning better estimates for uncertainty in ocean current predictions

Gaussian Processes: estimate the value **with an uncertainty estimate!**

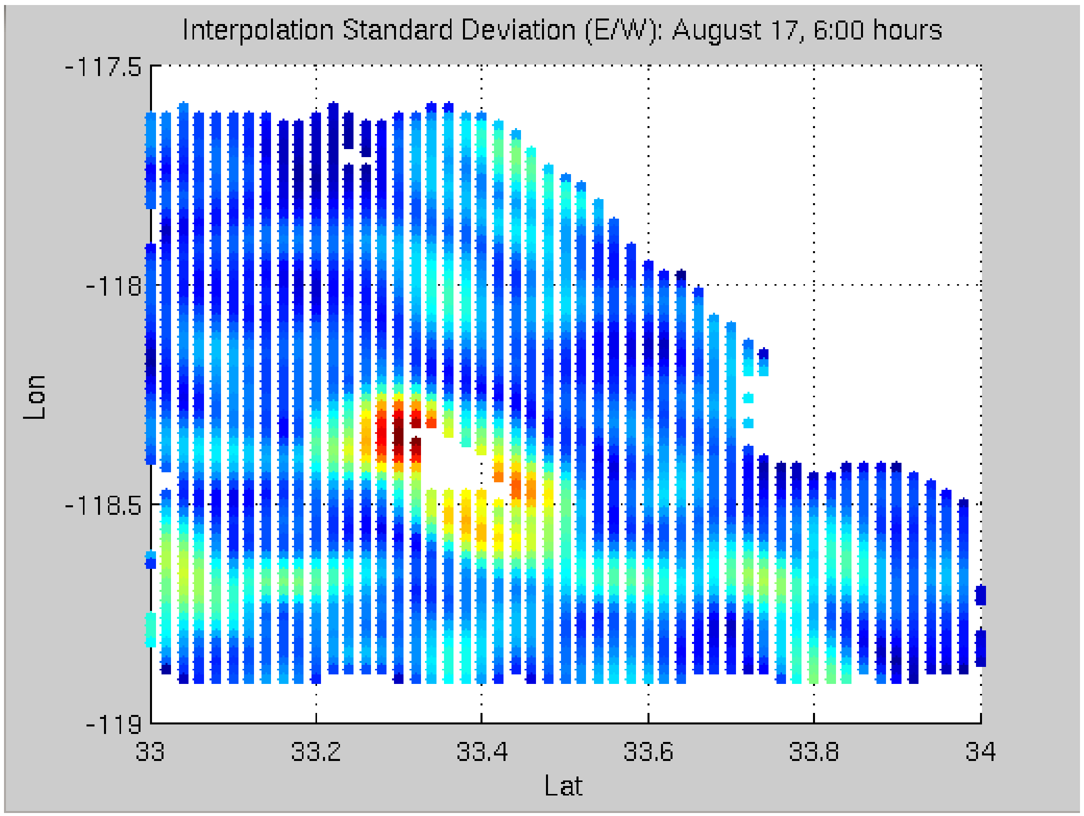

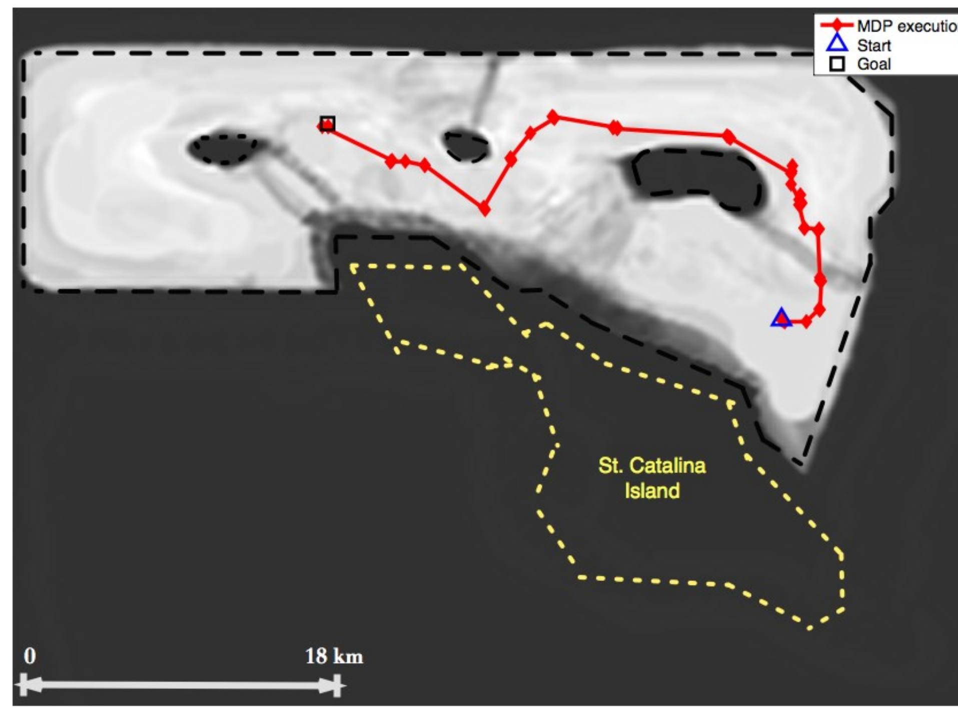

Planner	Noise	Pros	Cons
Minimum-Expected-Risk	Low variability currents	+ Goal-directed + Fast	- Poor in strong currents
Stationary Finite Horizon MDP	Low variability currents	+ Trade-off between goal-directed and risky behavior + Reasonably fast	- Stationarity assumption may be limiting
Non-stationary finite horizon MDP	High variability currents	+ Can take advantage of currents to cross risky sections	- Susceptible to timeouts - Computationally Expensive

Field testing!

Year	Planner	Field hours glider
2011	Min-Risk	408
2012	Stationary MDPs without GP predictions	840
2012	Minimum-Expected-Risk planner	360
2012	Stationary GP-MDP	120
2013	Non-Stationary GP-MDP	168
2011-13	**Total**	**1896**


GARMIN
UNIVERSITY OF
SOUTHERN CALIFORNIA
RUSALKA
Trojans
SC
STOP the
FOULING!

Can we develop systems that utilize the currents?

Can we develop systems that utilize the currents?

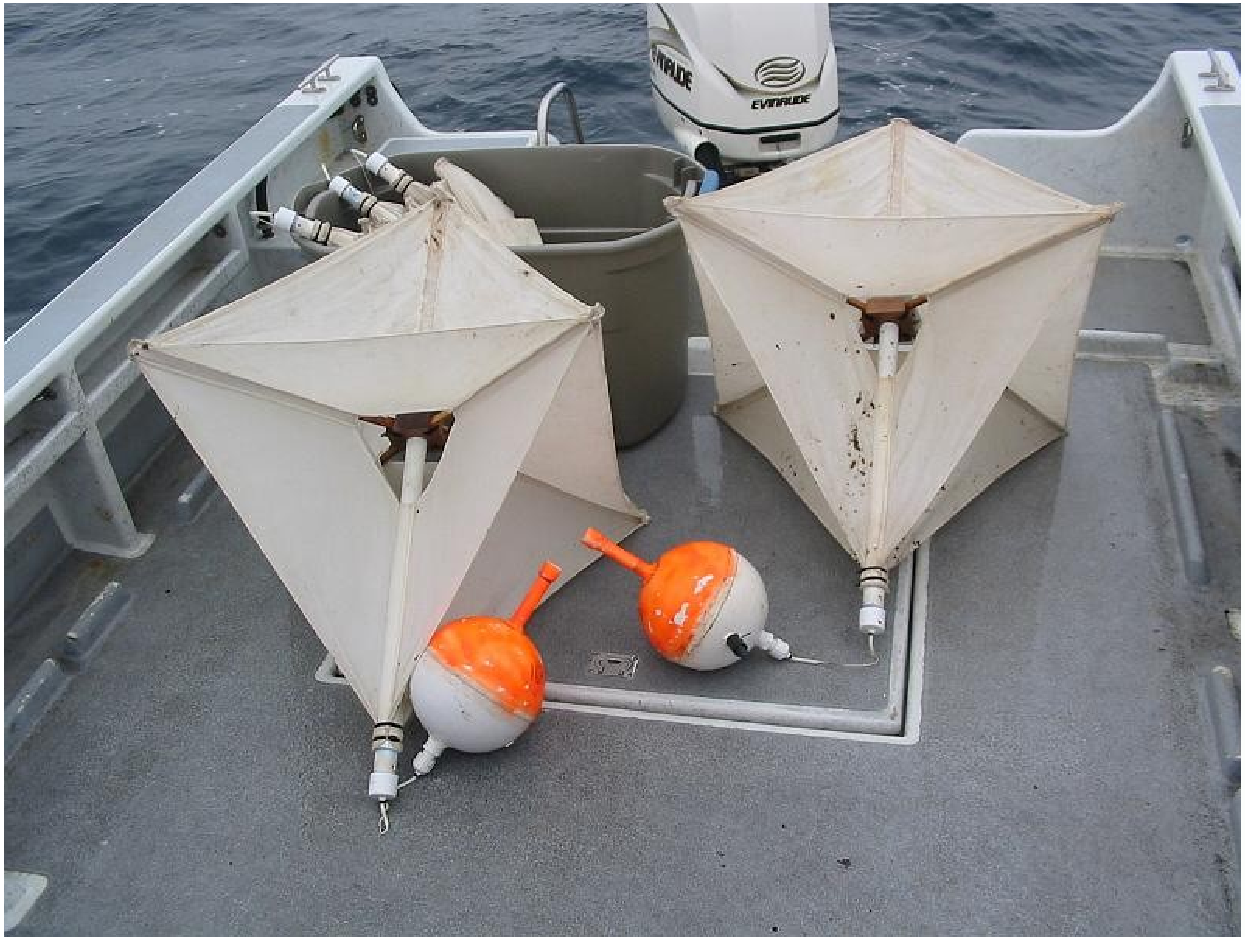

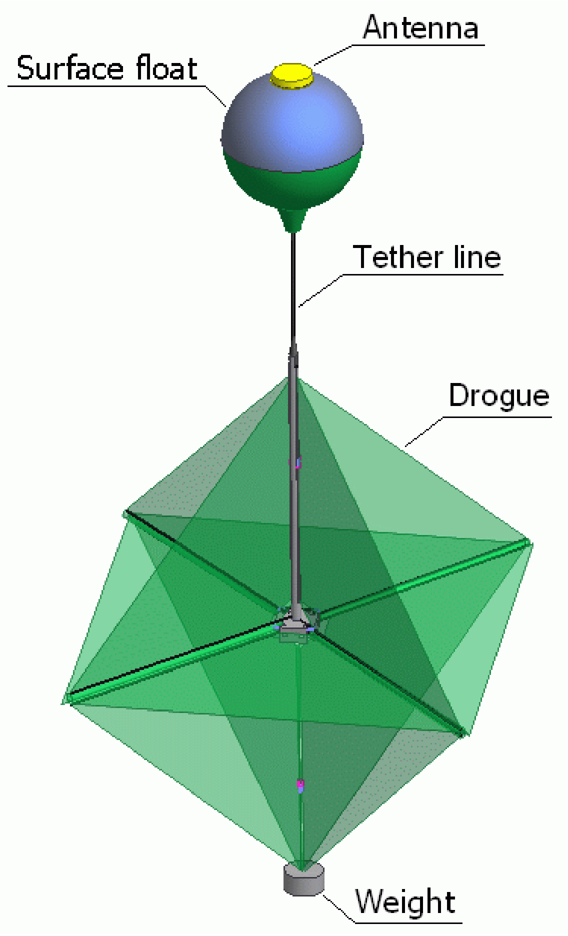

Microstar drifter, Pacific Gyre

Active drifters

“Choose the current to take you where you want to go”

Added benefits:

Easy deployment

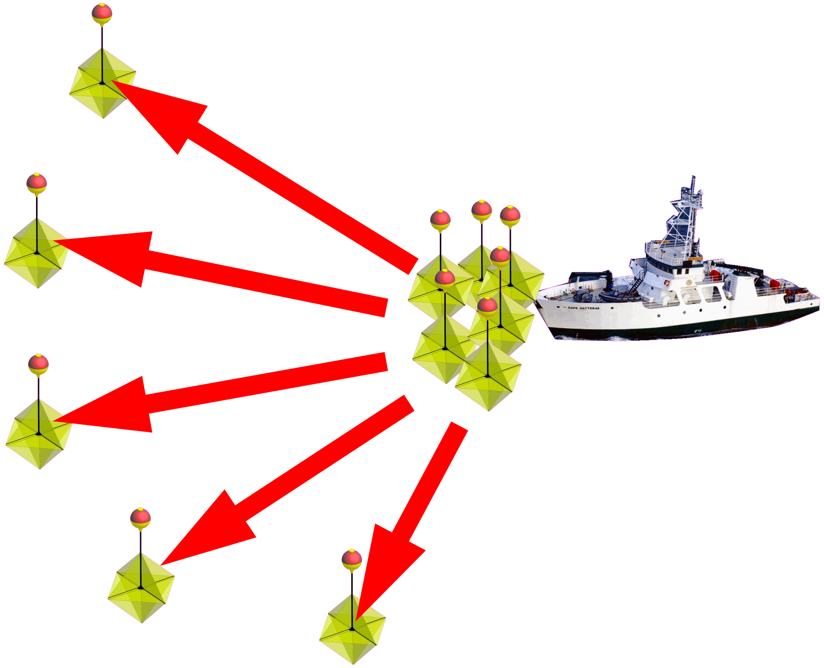

Efficient recovery

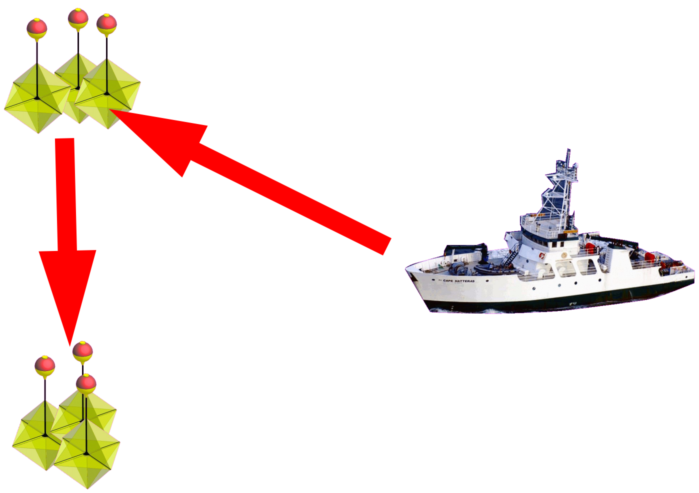

Simulation experiments using ROMS

When to pick a new current ?

- Track angle between desired direction of movement and the current movement

How to pick a new current?

- Pick depth where current at desired direction

How to coordinate?

- Closely located drifters can share current estimates

Simulation results: can collect drifters in few clusters.

Example deployment

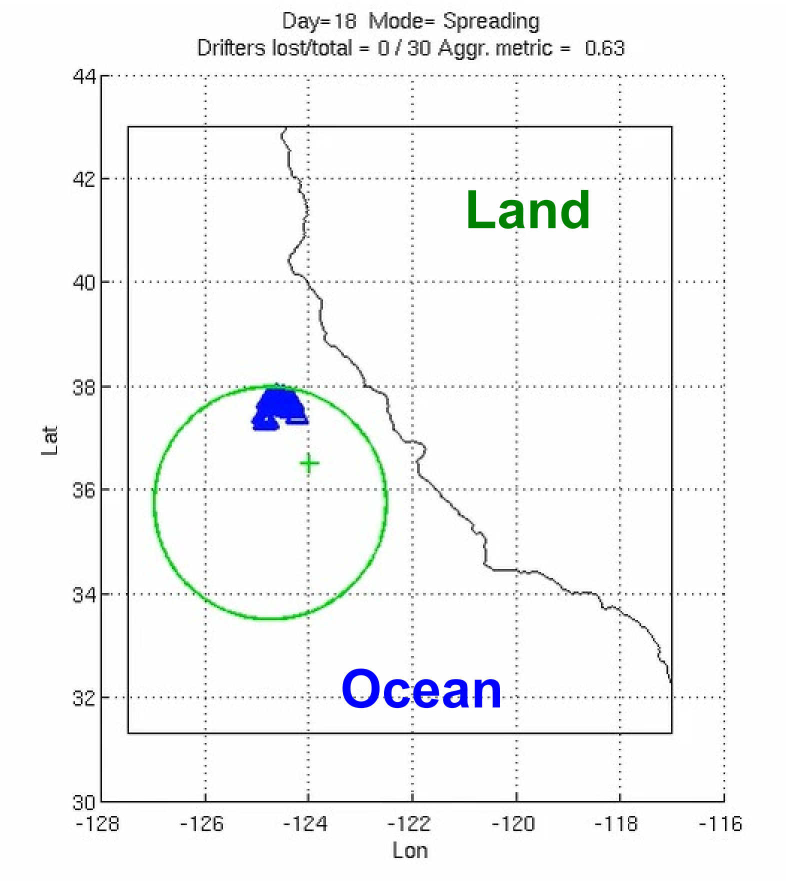

Aggregation performance over 100 simulations

	Average metric		
	90 days	180 days	
Ideal drifter	0.49	0.55	~ 2 clusters
Realistic drifter	0.21	0.29	~ 3-5 clusters

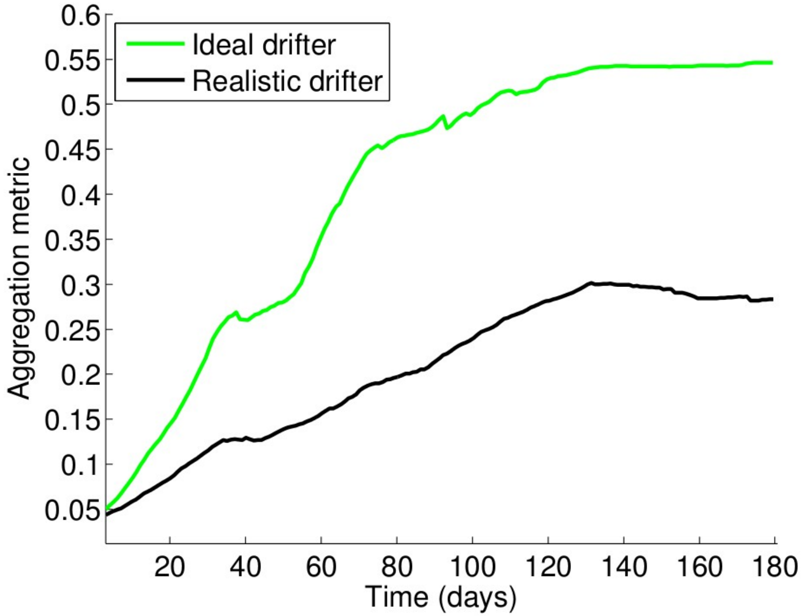

Artem Molchanov, Andreas Breitenmoser

What if there is no appropriate sensor, and the biology needs to be analyzed in the lab?

Ex-situ sampling

Lab analysis of physical samples, labeled offline in batches

MBARI Dorado AUV
Ten 1.8 L gulpers
can fill once!

lab analysis

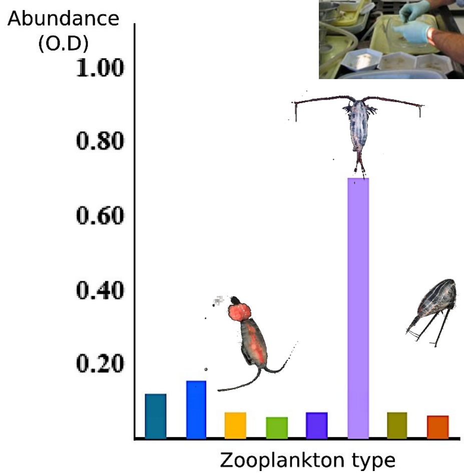

organism abundance

Given a limited number of gulpers, when to sample?

Learn from previous data when to sample

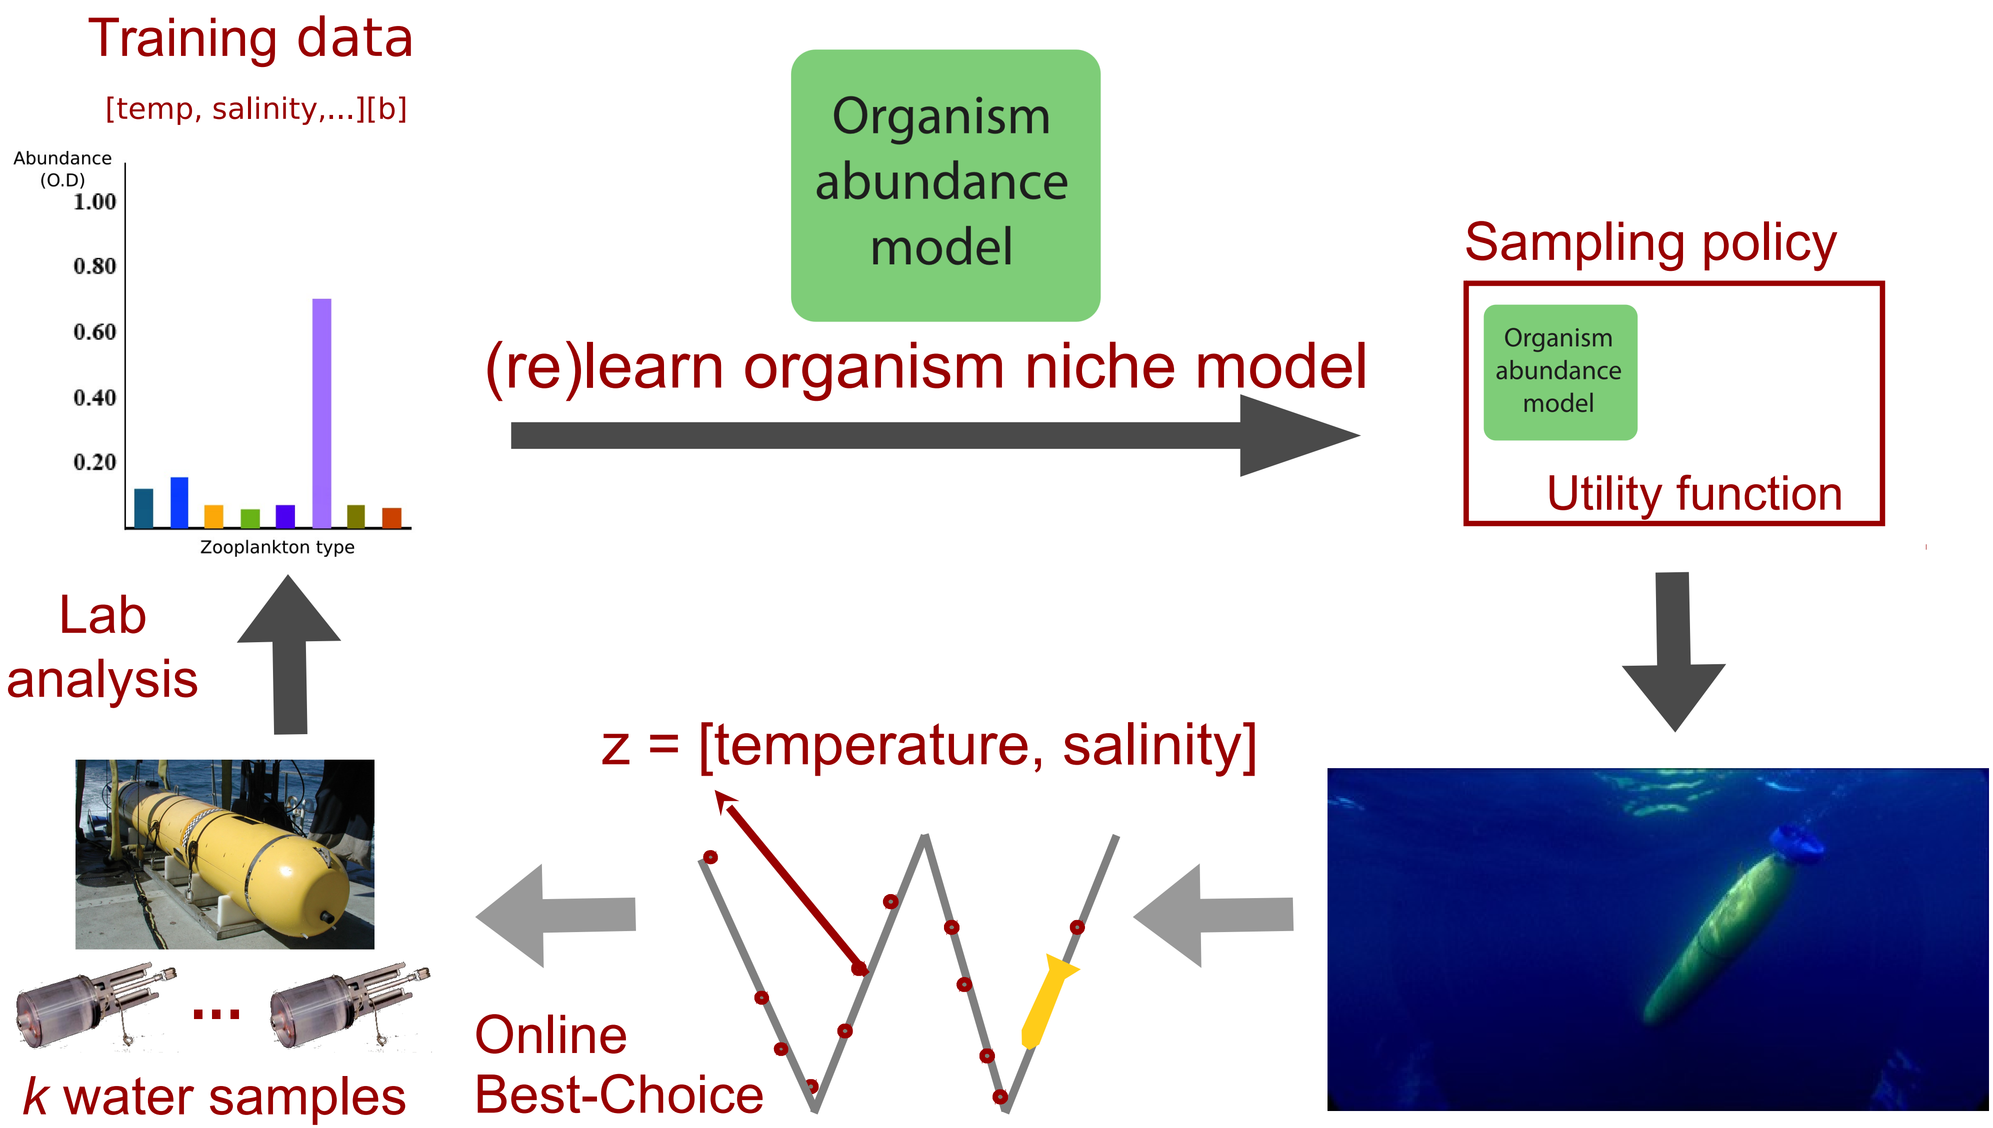

Online best-choice problem

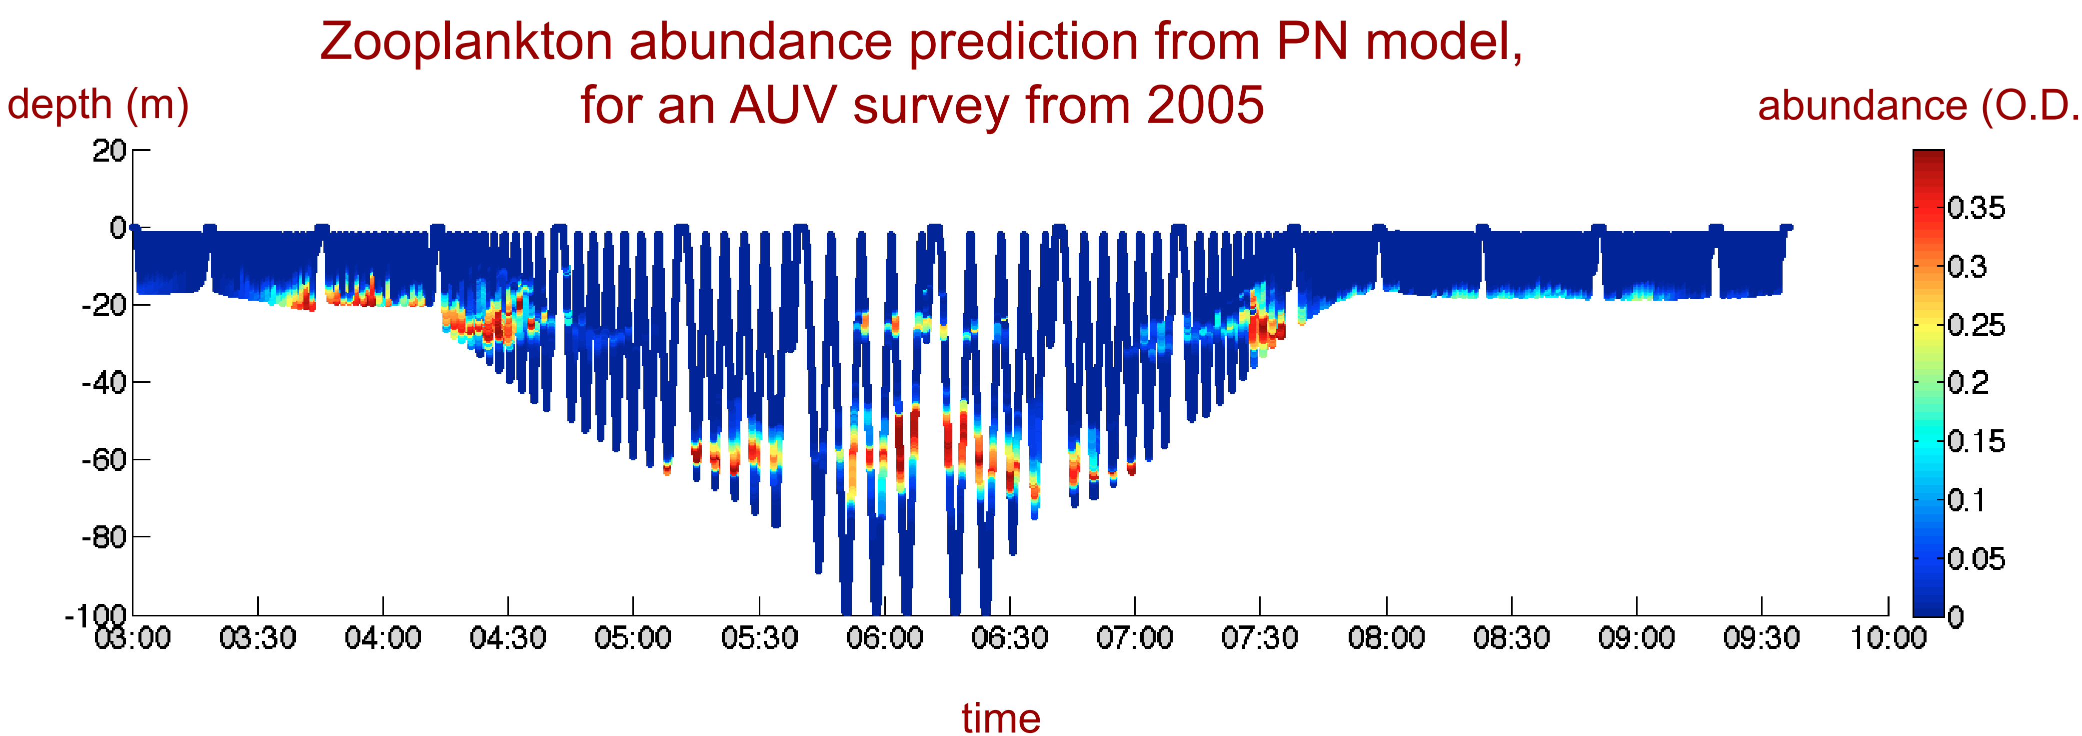

How to choose *k* samples to maximize the sum of utility from all samples?

Optimal Stopping Theory

Choose when to take a particular action.

The Hiring Problem:

- N candidates arrive for an interview i.i.d, and ranked
- Goal: choose single best candidate, in an online fashion
- Hiring decision is irrevocable! → can only gulp once!

Optimal Stopping Theory

Choose when to take a particular action.

The Hiring Problem:

- N candidates arrive for an interview i.i.d, and ranked
- Goal: choose single best candidate, in an online fashion
- Hiring decision is irrevocable! → can only gulp once!

Solution:

- Observe first N/e (36.7 %) candidates, then hire next best
- If there is no better candidate, hire the last person
- Guarantee: Probability choosing best candidate = 1/e (~36.7 %)

Selecting k candidates, online

Submodular hiring problem

- N candidates arrive for an interview, i.i.d, and rated
- Goal: choose best k candidates, online (best sum of rating)
- Hiring decisions are irrevocable → can only gulp once!

Selecting k candidates, online

Submodular hiring problem

- N candidates arrive for an interview, i.i.d, and rated
- Goal: choose best k candidates, online (best sum of rating)
- Hiring decisions are irrevocable → can only gulp once!

Solution

- Split total window into k segments
- Apply hiring algorithm in each segment
- Guaranteed competitive-ratio of at least (1 - 1/e)/11, ~0.05

Field trial

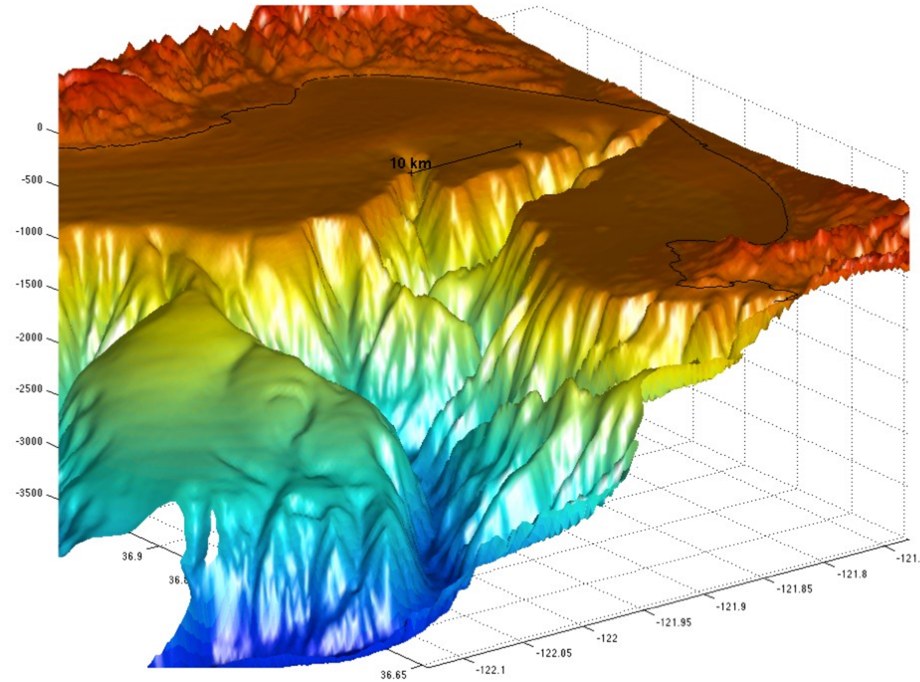

Dorado AUV on R/V Rachel Carson with the gulper bay open (Monterey Bay)

1 km x 1 km Lagrangian surveys depth ~30 m, duration ~4.5 hr

Field trial set-up

Goal : Acquire high abundance samples of pseudo-nitzschia (PN), a potentially toxinogenic alga

87 analyzed samples from October 2010 CANON experiment used to learn niche model for pseudo-nitzschia

Cross-validation to pick input variables and GP kernel parameter

Mission in North Monterey Bay to acquire 9 samples (1 gulper was non-functional)

Predictions of trained pseudo-nizschia model

Prediction (mean)

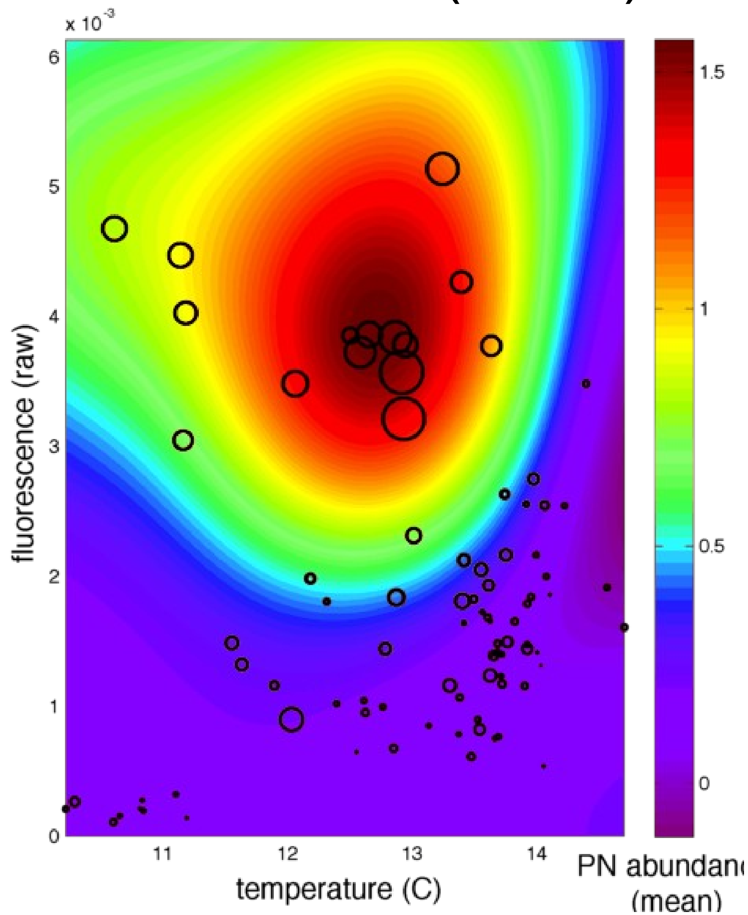

Uncertainty (variance)

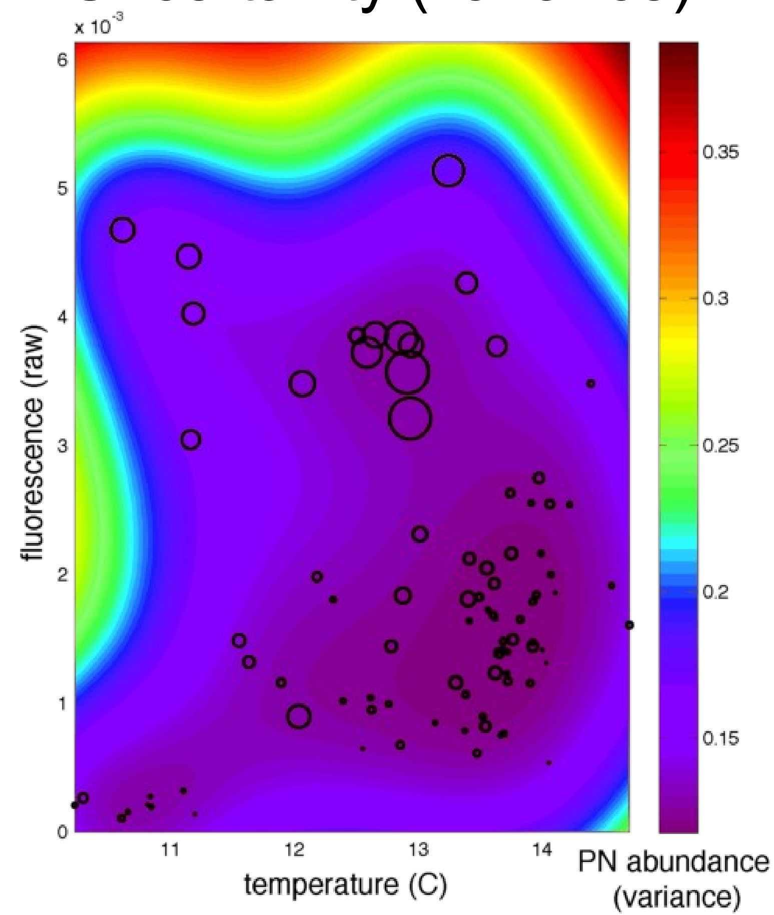

* circle size proportional to measured abundance

Samples acquired

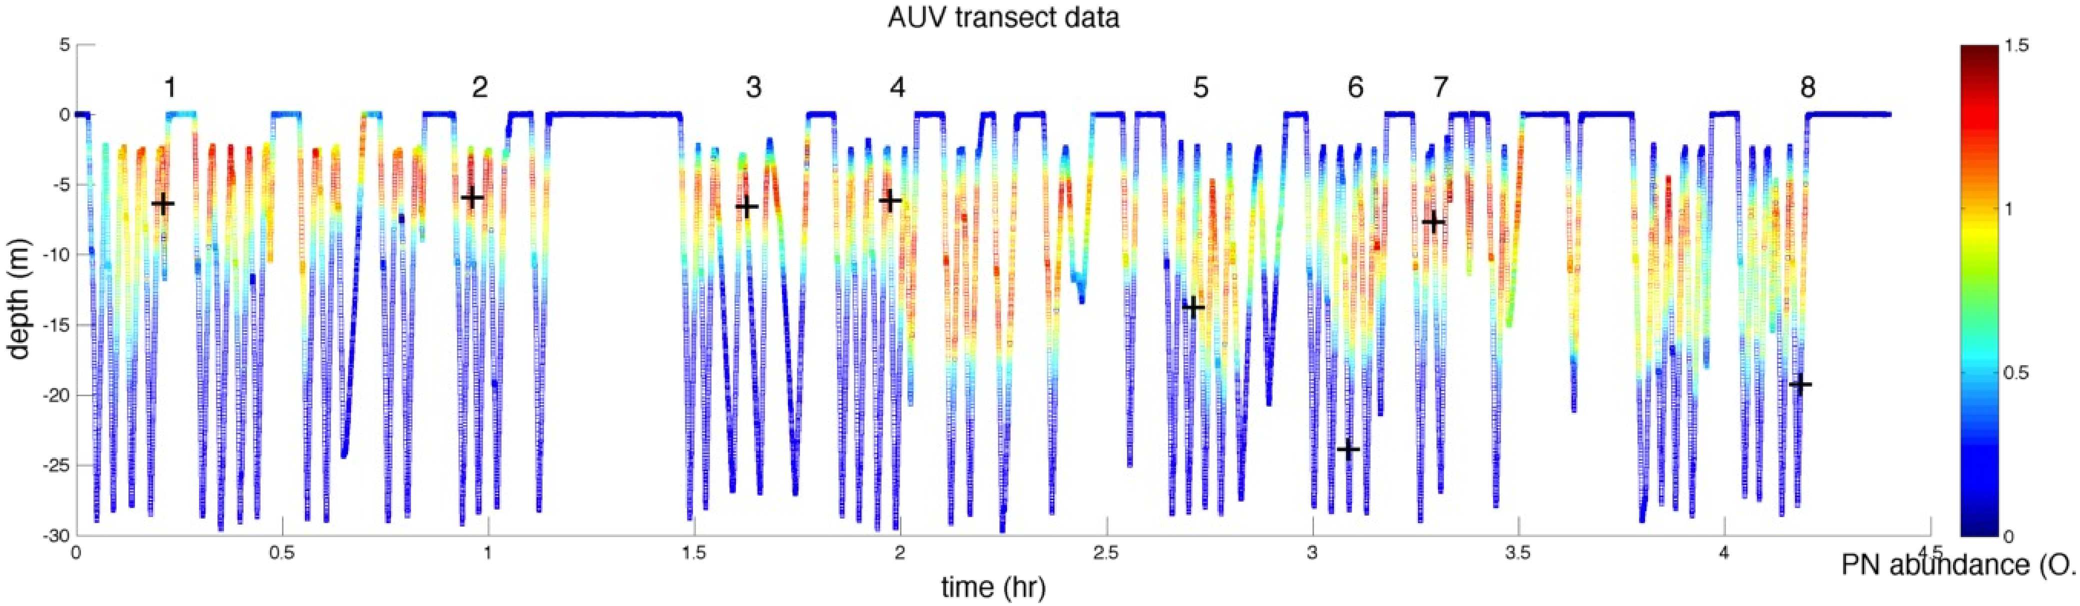

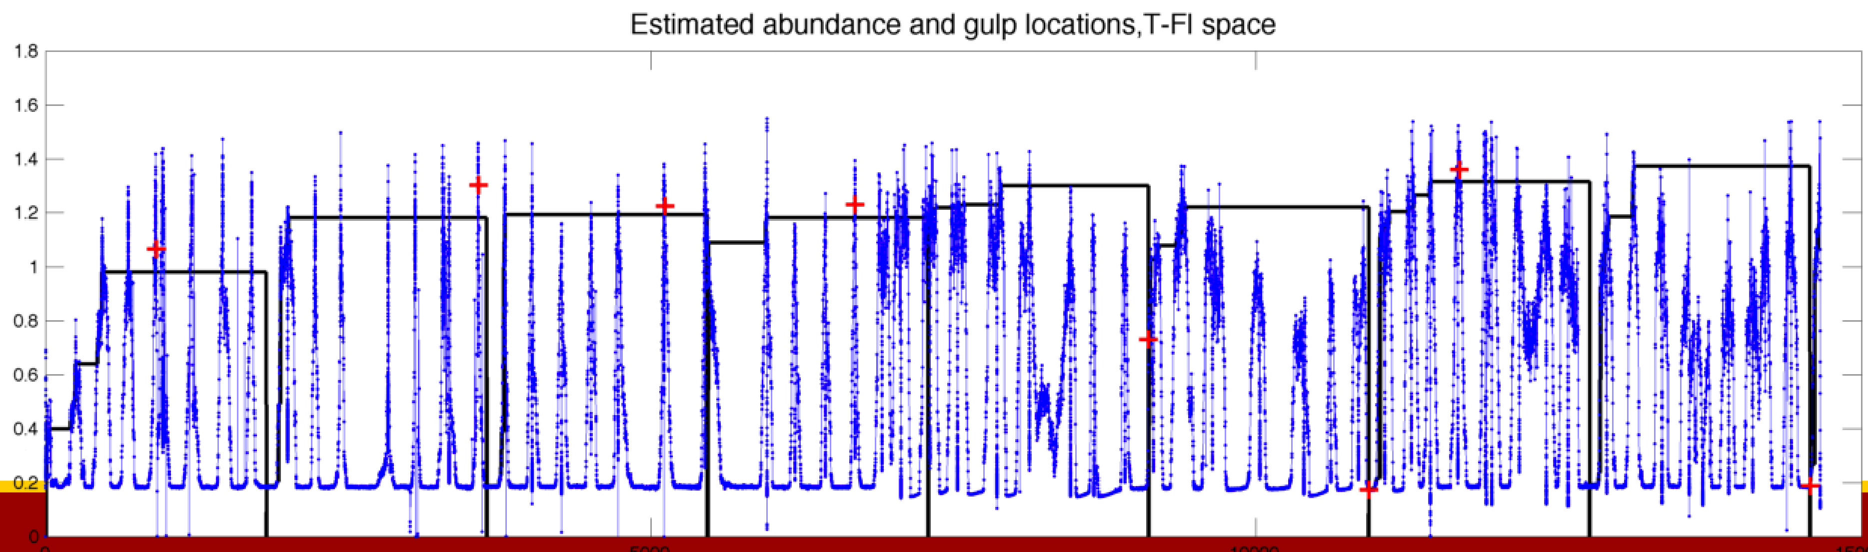

Ex-situ sampling contributions

- Stochastic, online constrained sampling
- Model is geography agnostic
- Closes autonomy loop on ecosystem monitoring – first data-driven experiment of this type in marine robotics
- Allow domain experts to task vehicles at a high(er) level ("bring me the harmful microbe!")

In-situ adaptive sampling

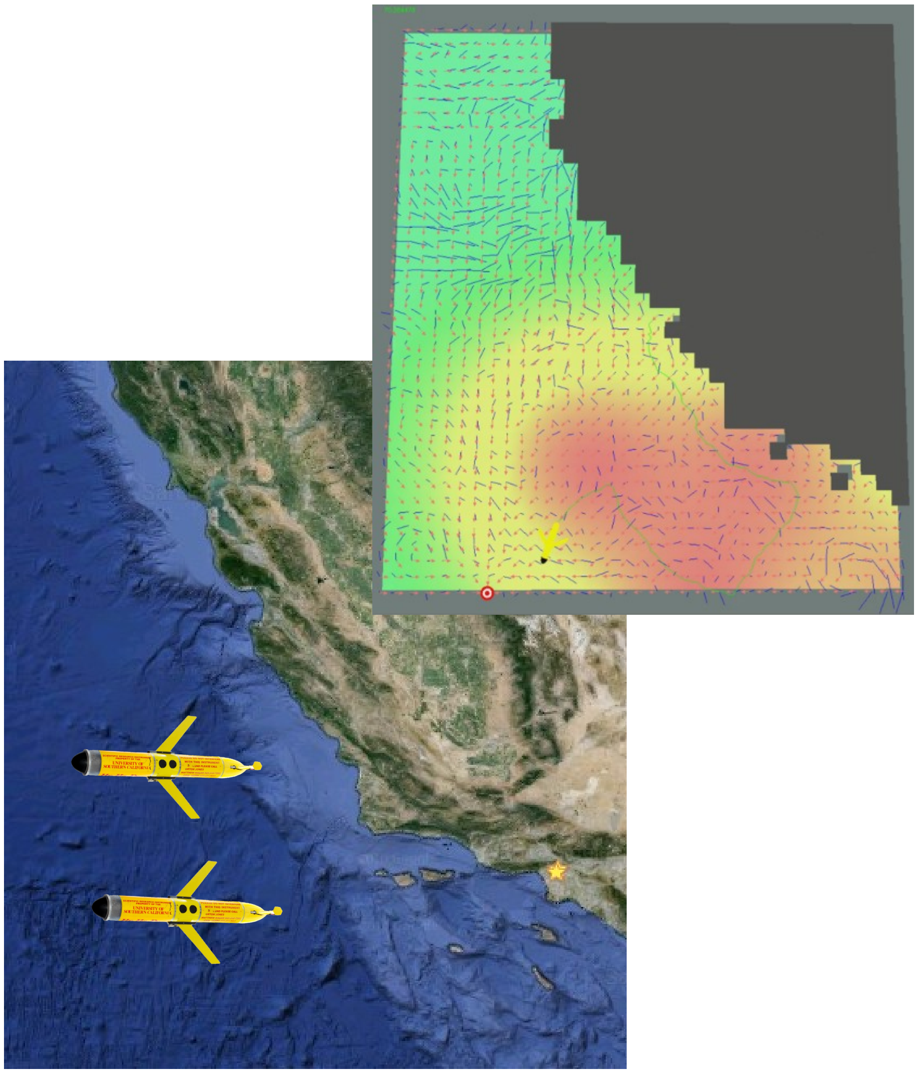

Lantao Liu, Kai-Chieh Ma, Stephanie Kemna

Online, adaptive sampling

- Adapt the vehicle movements based on its measurements, as the vehicle is sampling
- Create/update a model of the environmental phenomena

Informative Path Planning

- Gather the most informative data:
  Adaptive sampling using information-theoretic optimization criteria, such as entropy or mutual information
- Create the best model

Gaussian Process Regression Intro

A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution.

A Gaussian process is completely specified by its mean function and co-variance function. We define mean function $m(\mathbf{x})$ and the covariance function $k(\mathbf{x}, \mathbf{x}')$ of a real process $f(\mathbf{x})$ as

$$\begin{aligned} m(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})], \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))], \end{aligned} \qquad (2.13)$$

and will write the Gaussian process as

$$f(\mathbf{x}) \sim \mathcal{GP}\big(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')\big). \qquad (2.14)$$

[Rasmussen & Williams, 2006]

GP model selection

- choice of covariance function/kernel
    - common choice: squared exponential
- choice of hyperparameters
    - length scale
    - noise variance
    - signal variance

→ hyperparameter optimization, using prior data

GP prior & posterior

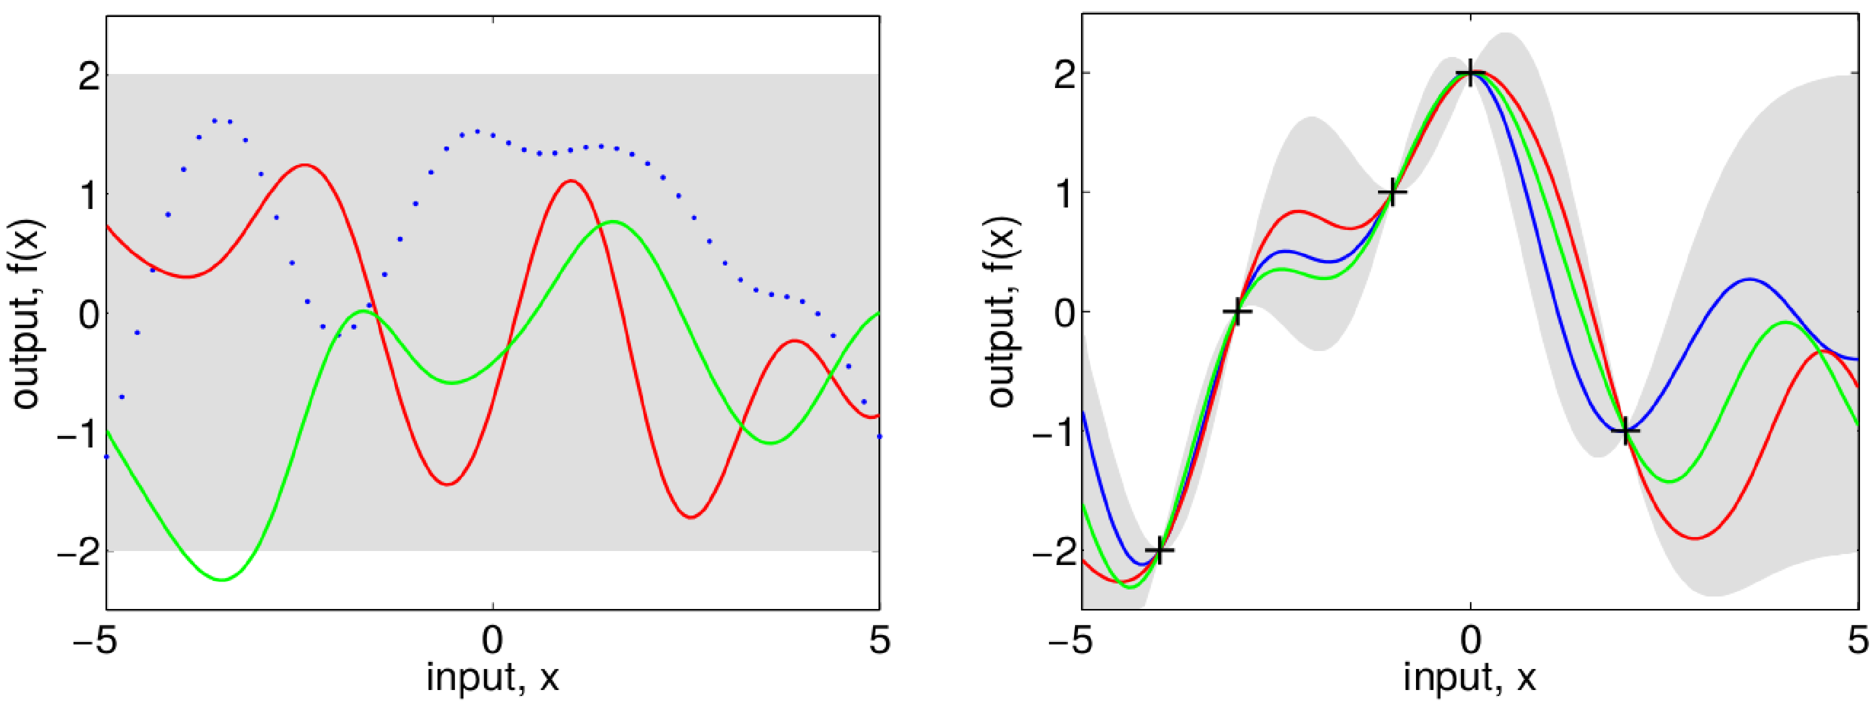

(a), prior (b), posterior

[Rasmussen & Williams, 2006]

Imagine; any location within your survey space can be represented by a Gaussian

Imagine; any location within your survey space can be represented by a Gaussian

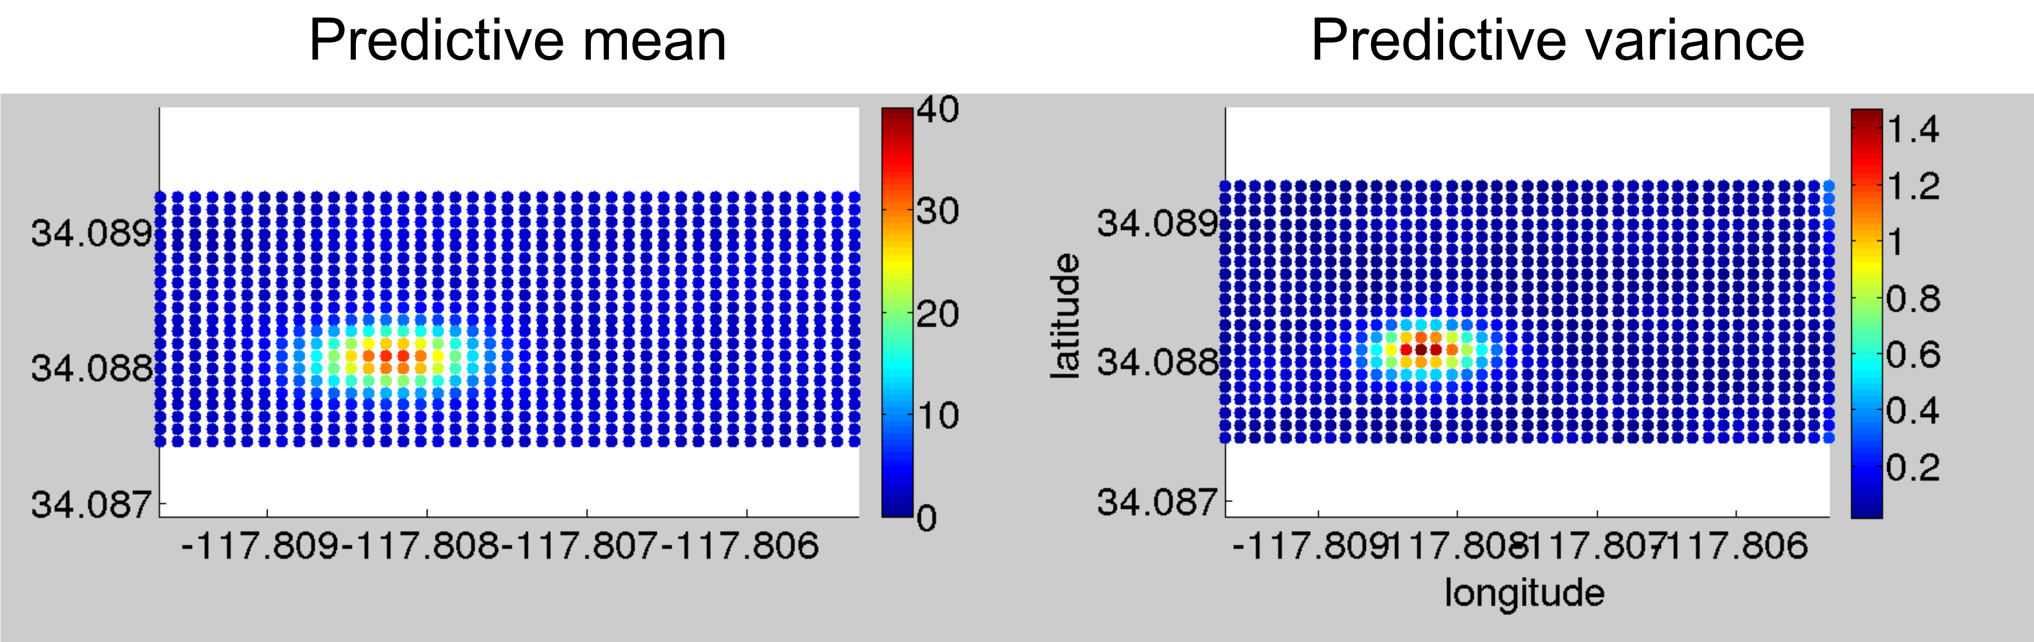

Metrics on GP output for determining quality of the environmental model

Quantify the uncertainty in the model, and calculate the information that can be gained for prospective sampling locations:

- Squared error
- Entropy
- Mutual Information
- Etc.

Path planning, given metric

- Greedy [Guestrin'05, Krause'08, Kemna]
    - local greedy [Low'12]
- Recursive Greedy; plan path from S to T [Binney'10, Krause'07, Singh'09]
- Dynamic Programming [Low'08/'09, Hitz'14, Ma/Liu]
- Branch & bound [Binney'12]

eMIP [Singh'06/'07/'09]

Informative path planning for AUVs

Ocean monitoring

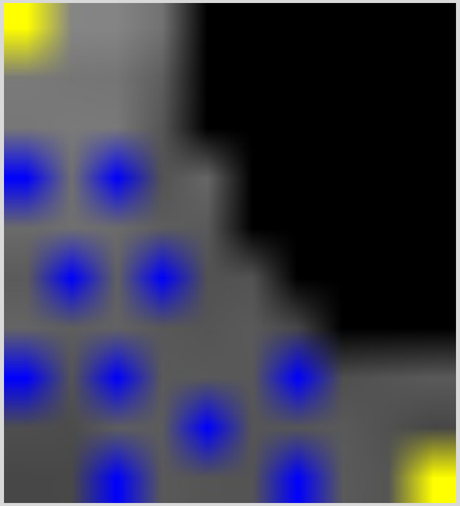

Potential sampling points

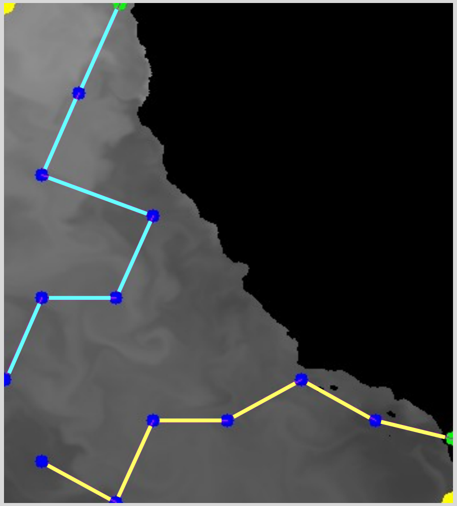

Planned paths

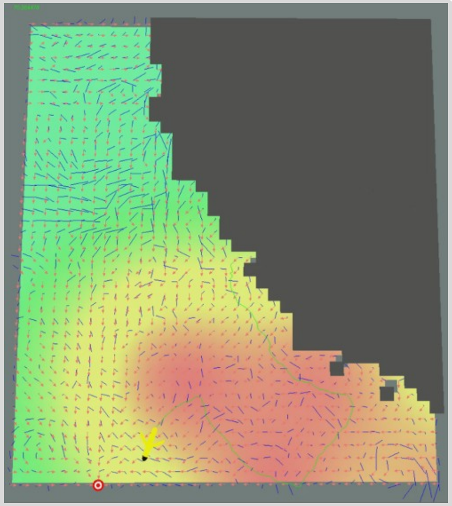

Informative sampling

Informative path planning for underwater glider – hierarchical planner

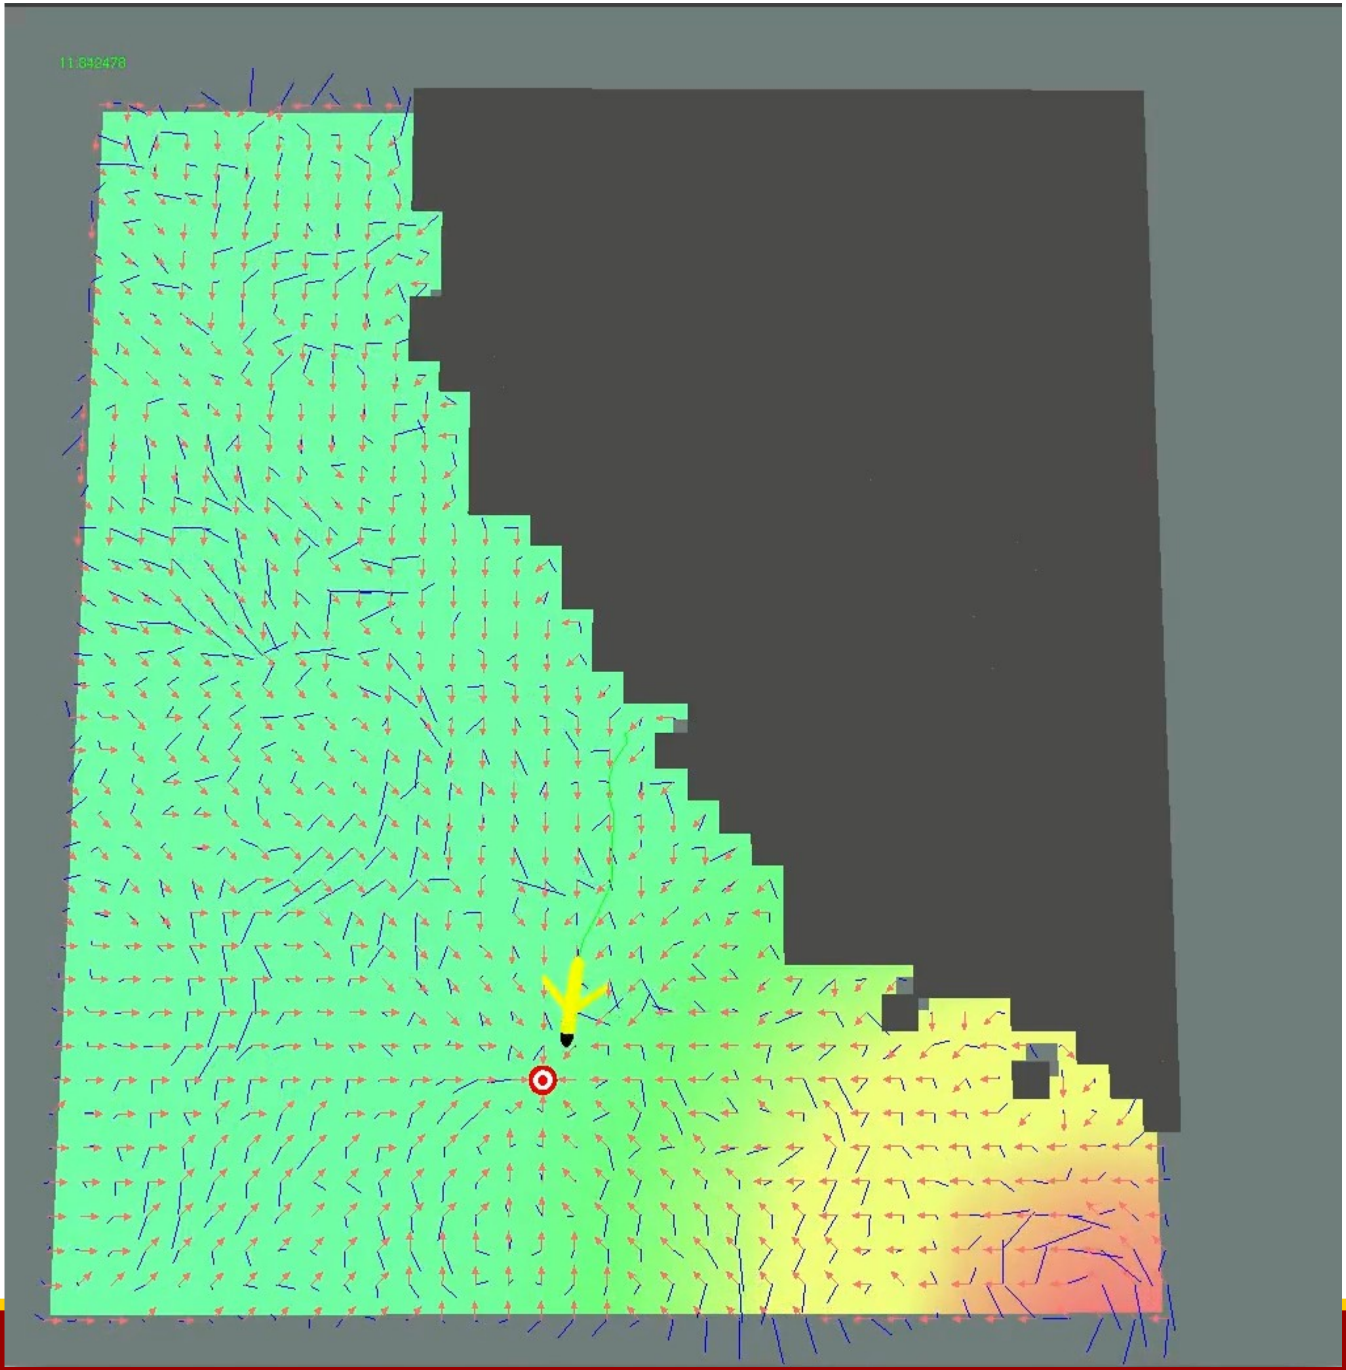

Lantao Liu, Kai-Chieh Ma

Adaptive versus standard surveys ?

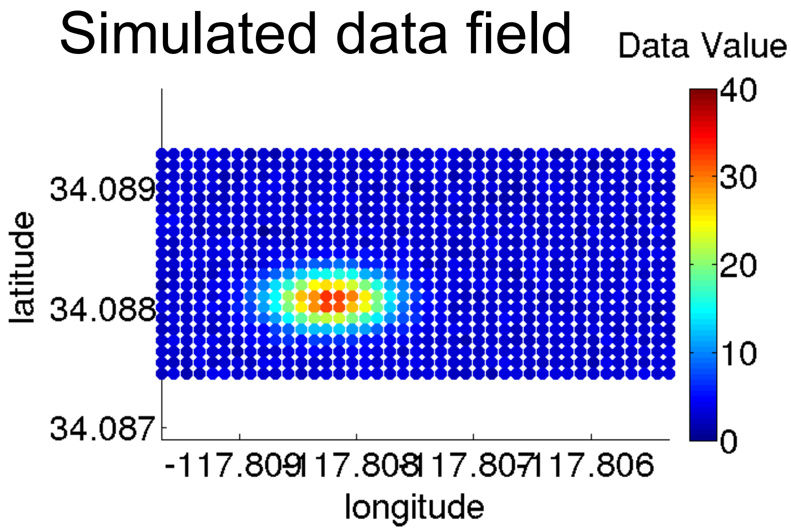

Choice of vehicle trajectories:

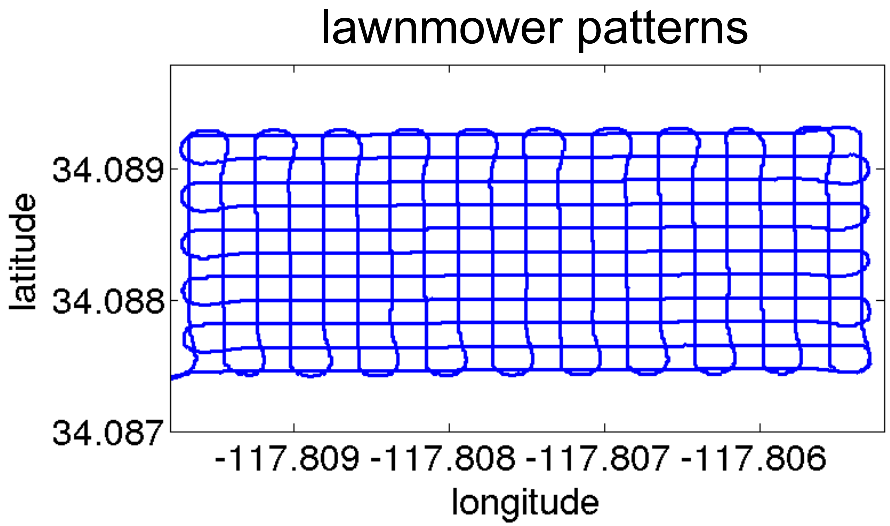

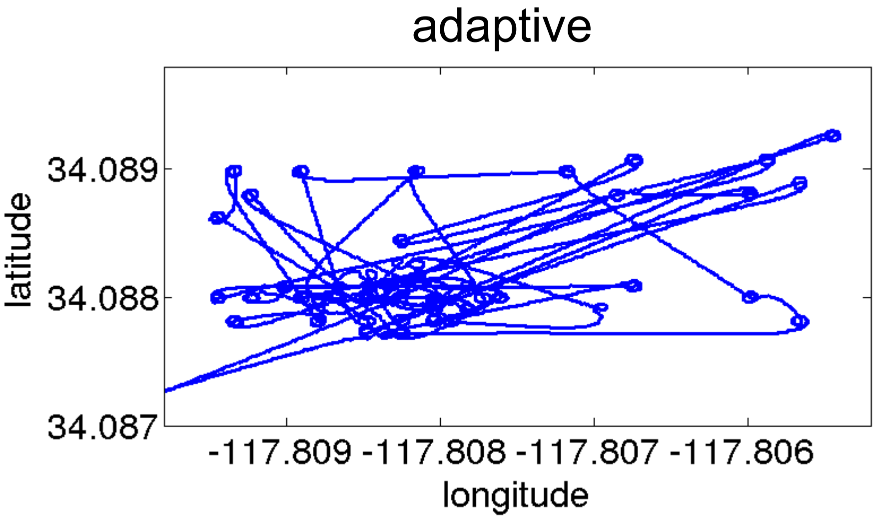

Benefits of informative path planning

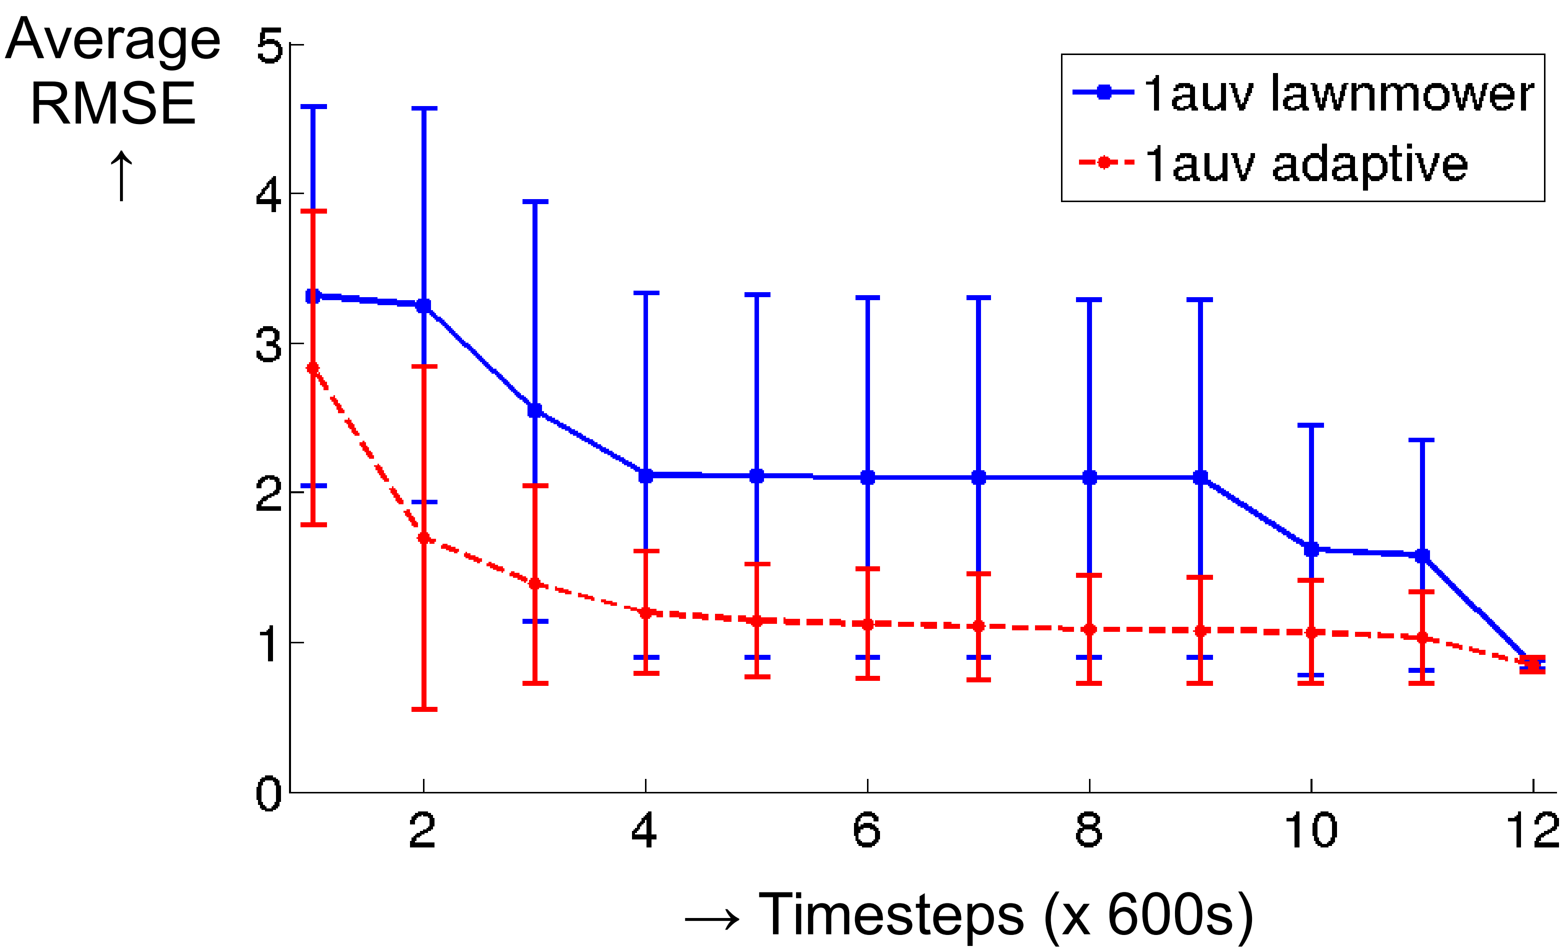

How to make sure the vehicle can operate safely in a previously unexplored environment?

Obstacle detection from overhead imagery using self-supervised learning

- Deploy robots in new environments with low risk
- Obstacle maps not available
- Need maps to plan paths
- Want to generate relevant maps without human labor

Combining aerial & sonar data

Aerial imagery

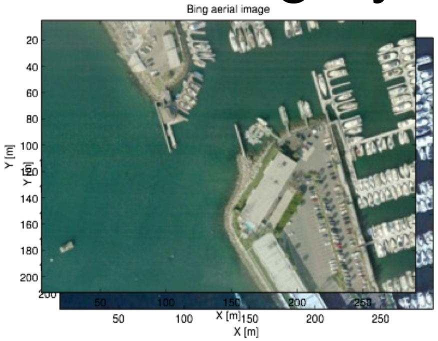

Feature extraction + Training labels generation

Sonar data

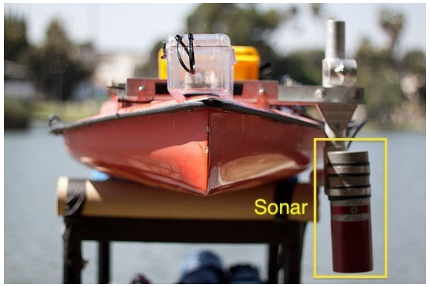

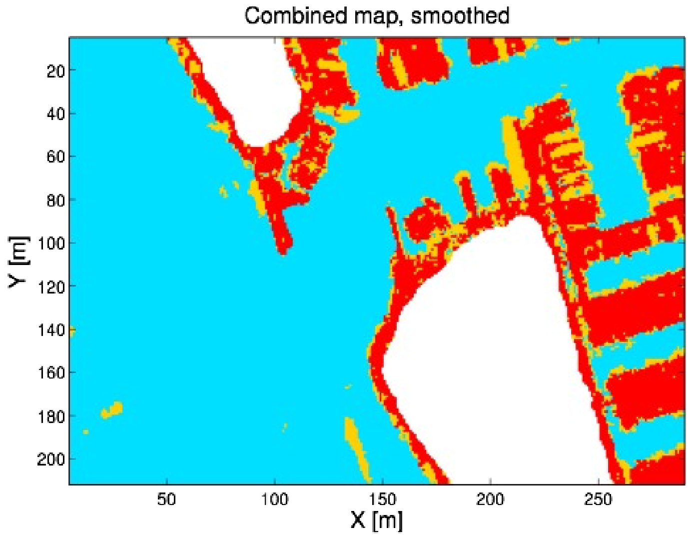

Prediction & smoothing

obstacle, transient, free space

What about in-field obstacle avoidance?

Different sensors for different parts of the environment

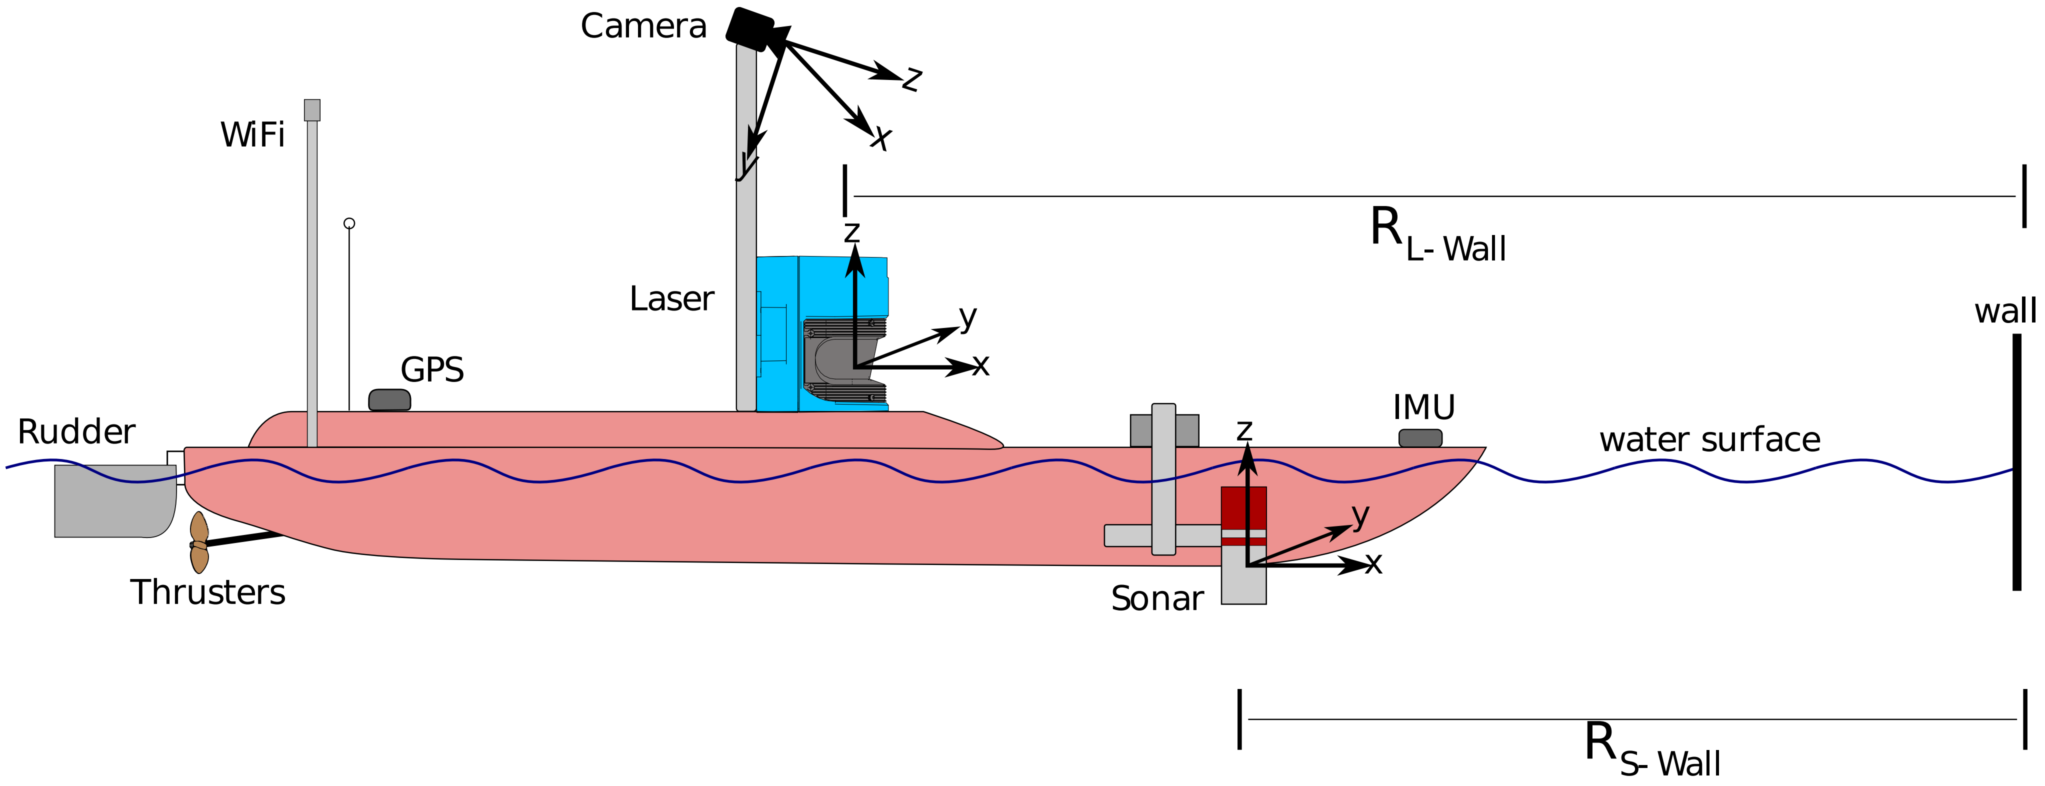

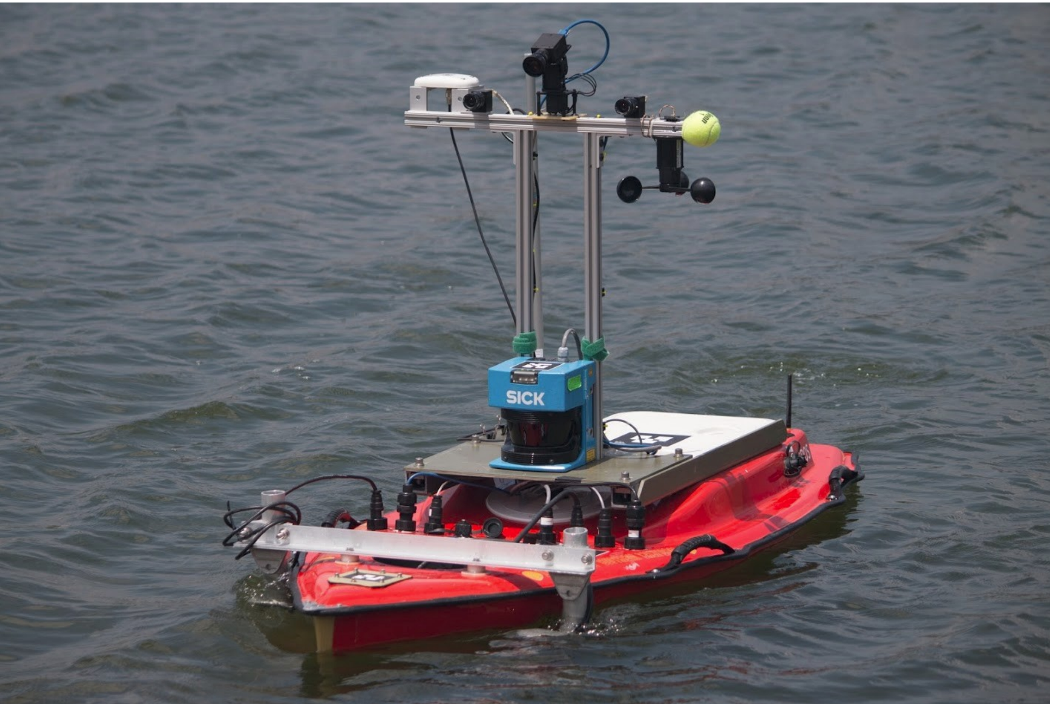

Hörður Heiðarsson
Stephanie Kemna

Autonomous sensor calibration

- Determine transformations between our different sensors:
  - Laser – Sonar:
    2D affine transform: translation, rotation, scaling
  - Camera - Water plane:
    6 DOF rigid body transform
- Actively gather data for calibration using existing features as calibration targets

Suitable calibration targets

- Sloped targets not suitable
- Straight edges give ambiguity
- Use corner features
    - Can be detected by our different sensors
    - Rarely sloped
    - Can be detected from overhead imagery

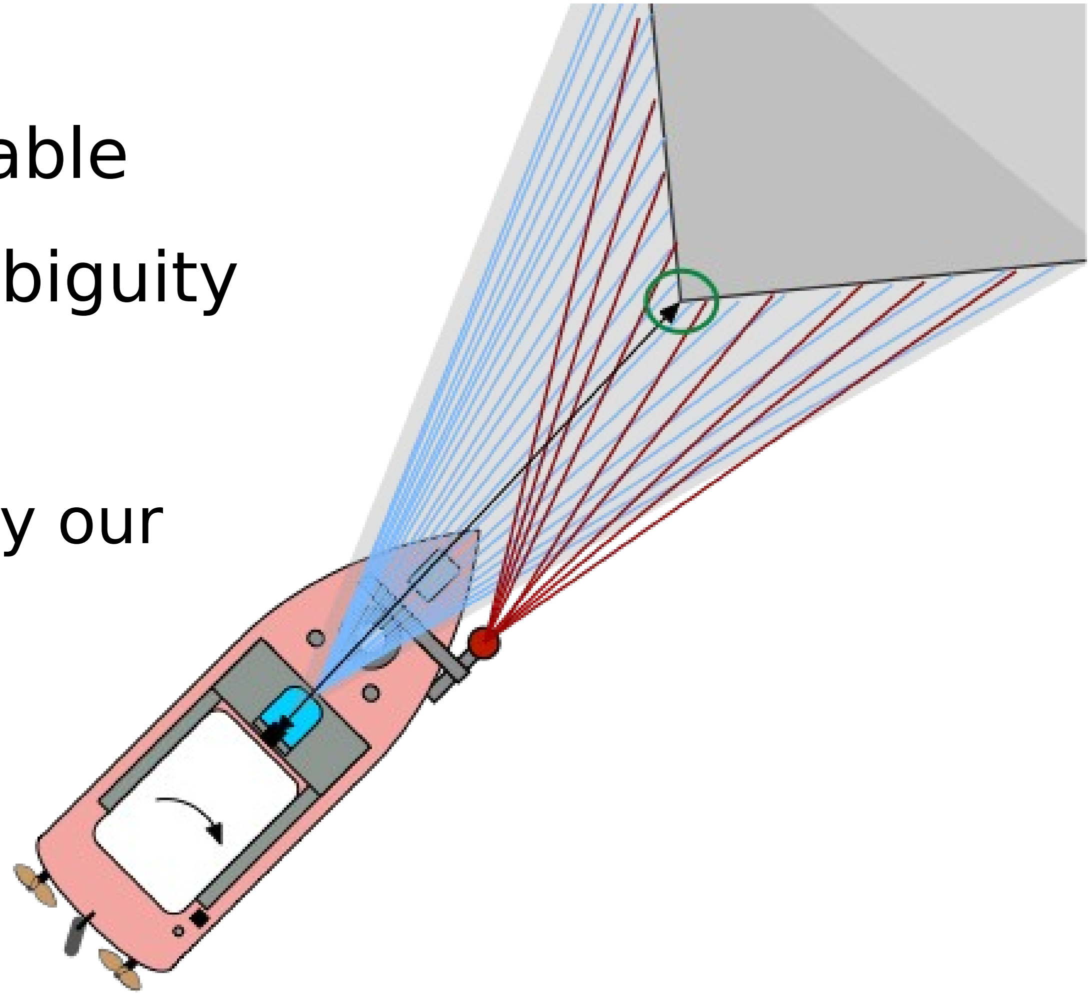

Feature extraction

For all sensors:

- Line extraction
- Find corners
- Run optimization to find best match between sensors

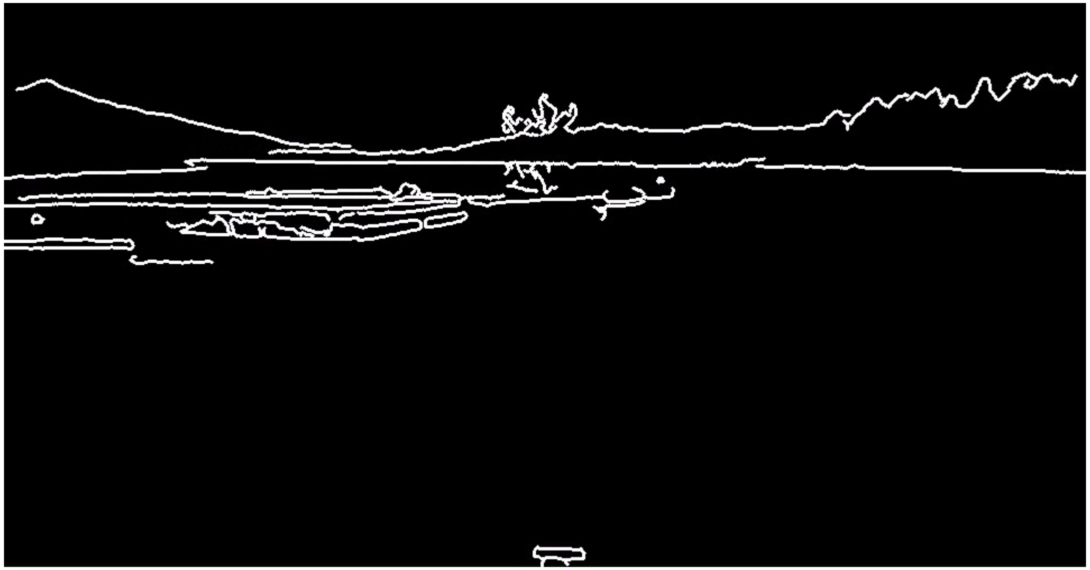

Results: laser & sonar

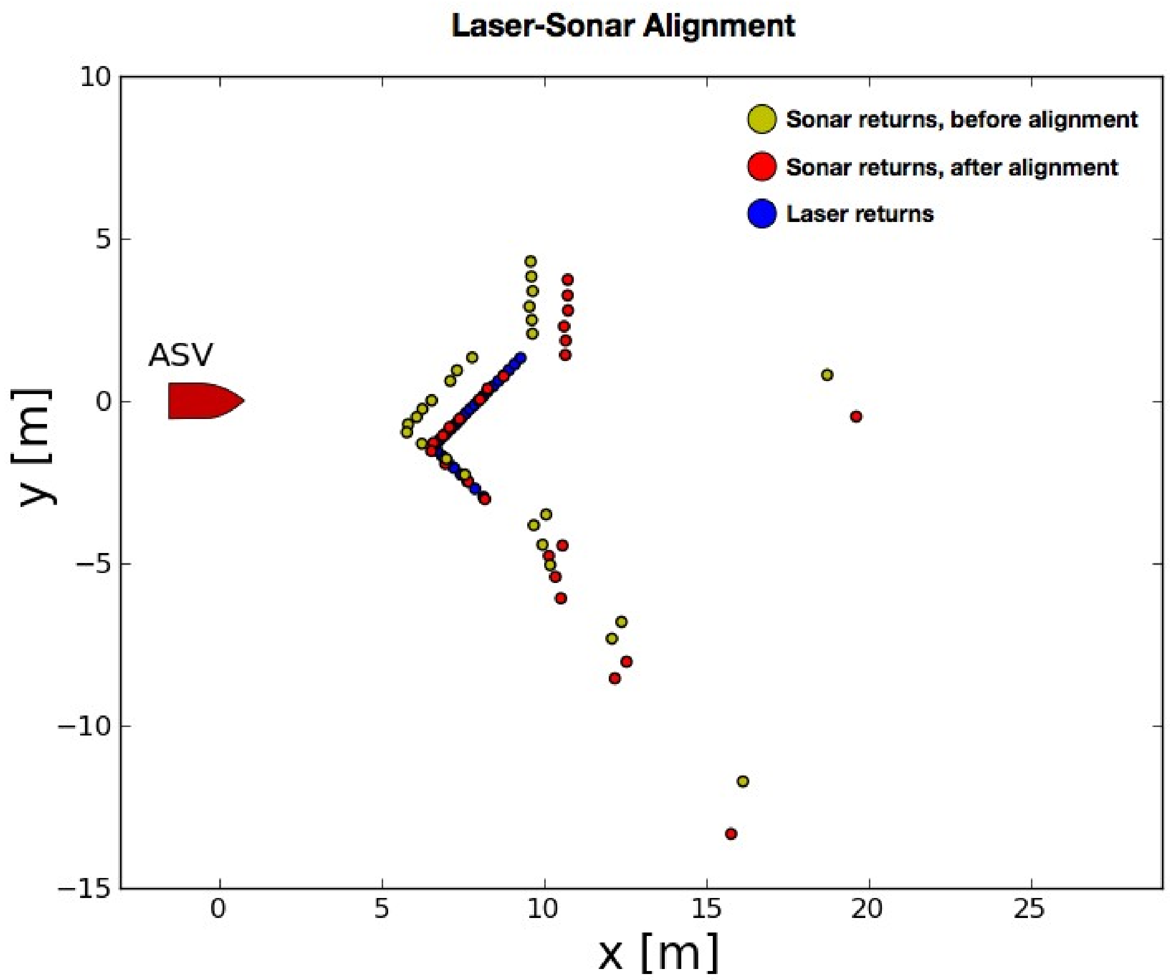

Multi-robot approaches

Multi-robot: run in parallel or coordinate?

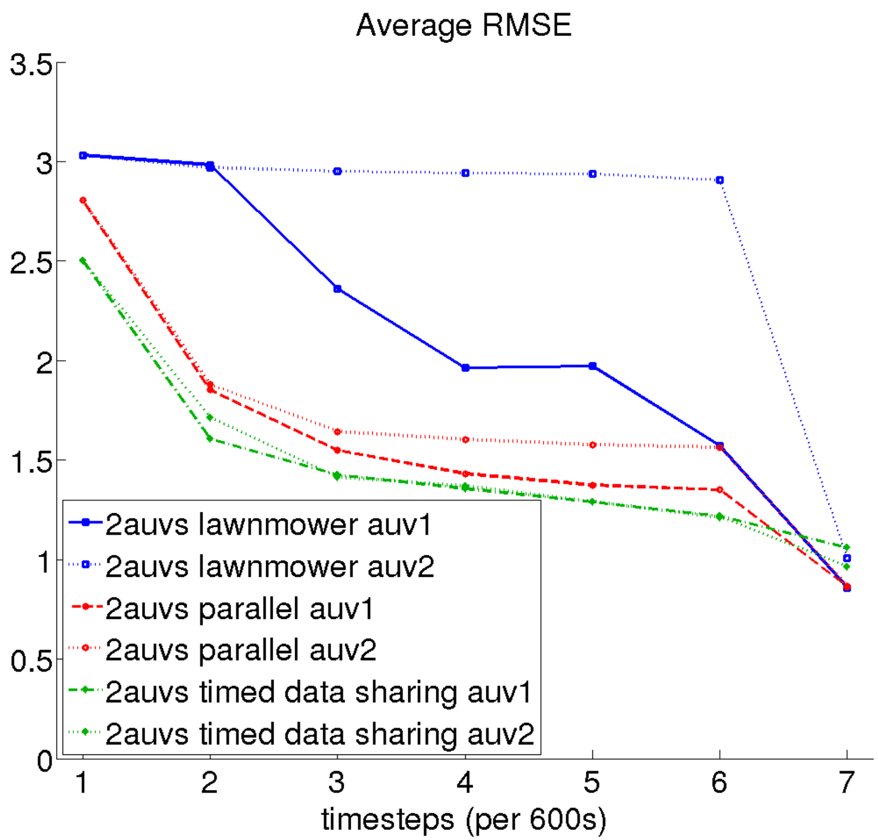

Orienteering solution from transformed matching graph

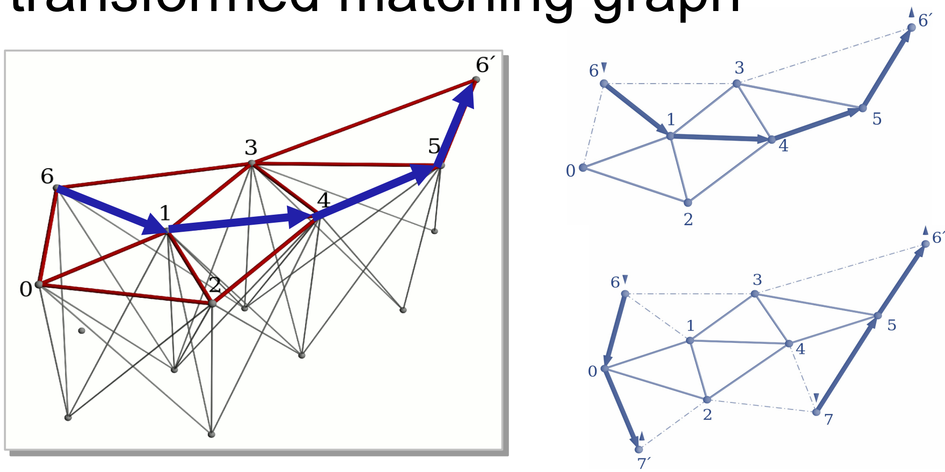

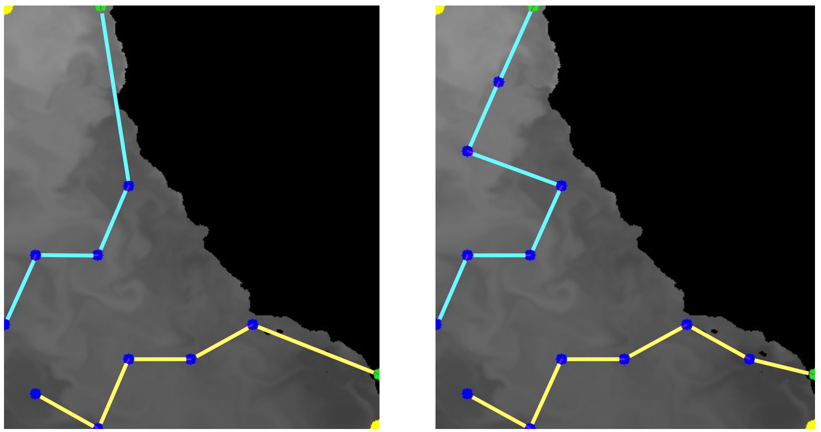

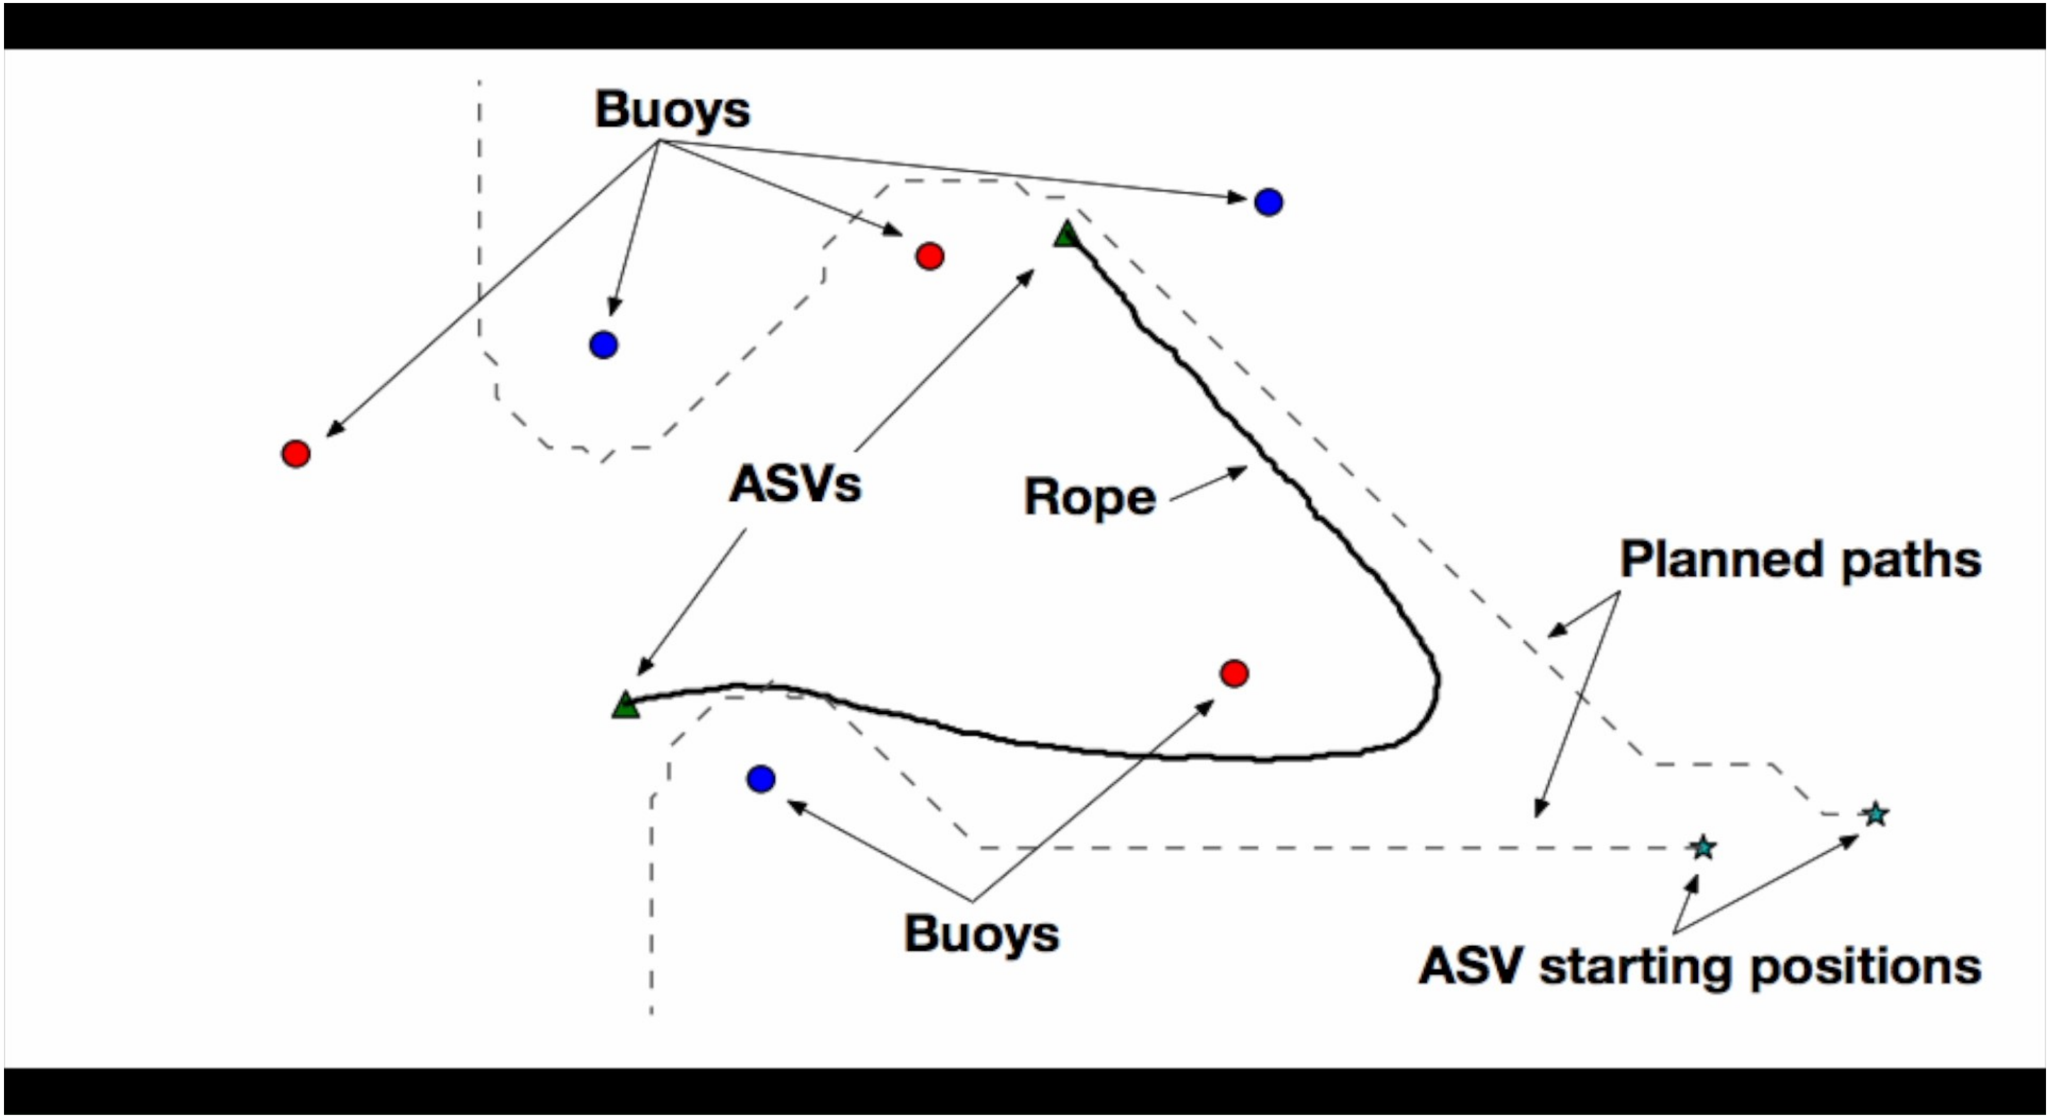
Buoys
ASVs
Rope
Planned paths
Buoys
ASV starting positions

What goes into getting overhead imagery at a lake…

Hordur Heidarsson, Jnaneshwar Das
Supreeth Subbaraya, Stephanie Kemna

What goes into getting overhead imagery at a lake...

Thank you!

http://robotics.usc.edu/resl/

Prof. Gaurav Sukhatme

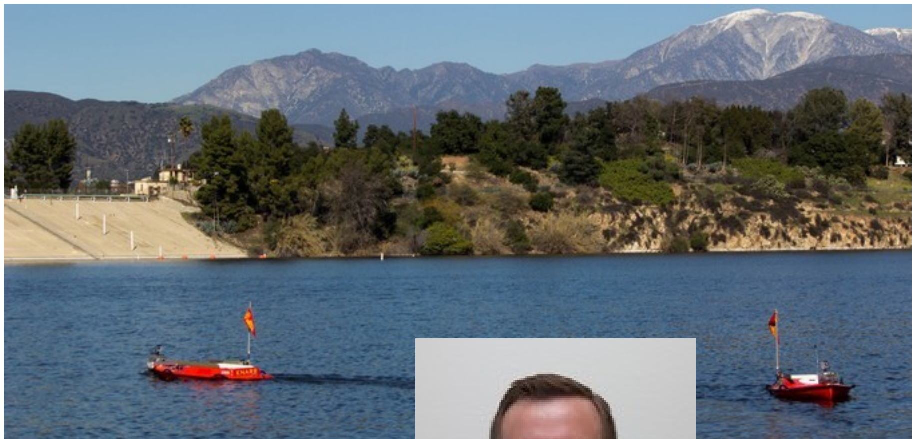

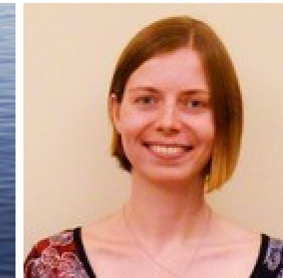

Stephanie Kemna

Andreas Breiten-moser

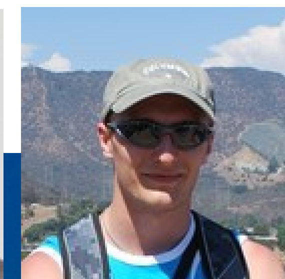

Artem Molchanov

Hörður Heiðarsson

Arvind Pereira

Jnaneshwar Das

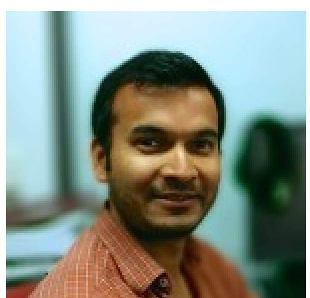

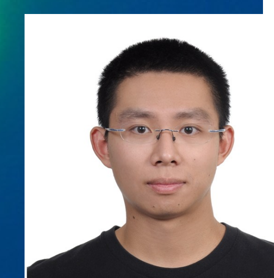

Kai-Chieh Ma

Geoff Hollinger

Lantao Liu

USC Viterbi School of Engineering

ROBOTIC EMBEDDED SYSTEMS LABORATORY

Planning to do aquatic robot experiments?

Remember to:

- always bring a towel
- use a canopy
- bring sunscreen & a cap
- bring an extra sweater, even in sunny SoCal!
- bring a rescue vehicle, e.g. kayak
- be prepared to talk football with the fishermen
- bring the internet

References

[Merckelbach, 2012] L. Merckelbach, “On the probability of underwater glider loss due to collision with a ship,” Journal of Marine Science and Technology, June 2012.

[Rasmussen & Williams, 2006] C. E. Rasmussen and C. K. I. Williams, Gaussian Processes for Machine Learning. The MIT Press, 2006.

Arvind Pereira, “Risk-aware path planning for autonomous underwater vehicles”, PhD Thesis, USC, 2013.

G. A. Hollinger, A. A. Pereira, J. Binney, T. Somers, G.S. Sukhatme, “Learning Uncertainty in Ocean Current Predictions for Safe and Reliable Navigation of Underwater Vehicles”, JFR 33(1), 2016.

A. Molchanov, A. Breitenmoser and G. S. Sukhatme. "Active Drifters: Towards a Practical Multi-Robot System for Ocean Monitoring". IROS, 2015.

Jnaneshwar Das, “Data-driven robotic sampling for marine ecosystem monitoring”, PhD Thesis, USC, 2014.

Hordur Heidarsson, “Obstacle Detection from Overhead Imagery using Self-Supervised Learning for Autonomous Surface Vehicles”, IROS, 2011.

Hordur Heidarsson, “Active Online Calibration of Multiple Sensors for Autonomous Surface Vessels”, ISER, 2014.

K. Ma, L. Liu, G. S. Sukhatme, “A Hierarchical Informative Path Planning Method for Ocean Monitoring.”, SCR, 2016.

S. Kim, S. Bhattacharya, H. Heidarsson, G. Sukhatme, V. Kumar, “A Topological Approach to Using Cables to Separate and manipulate Sets of Objects”, RSS 2013, IJRR 2015.

More publications are on our website!

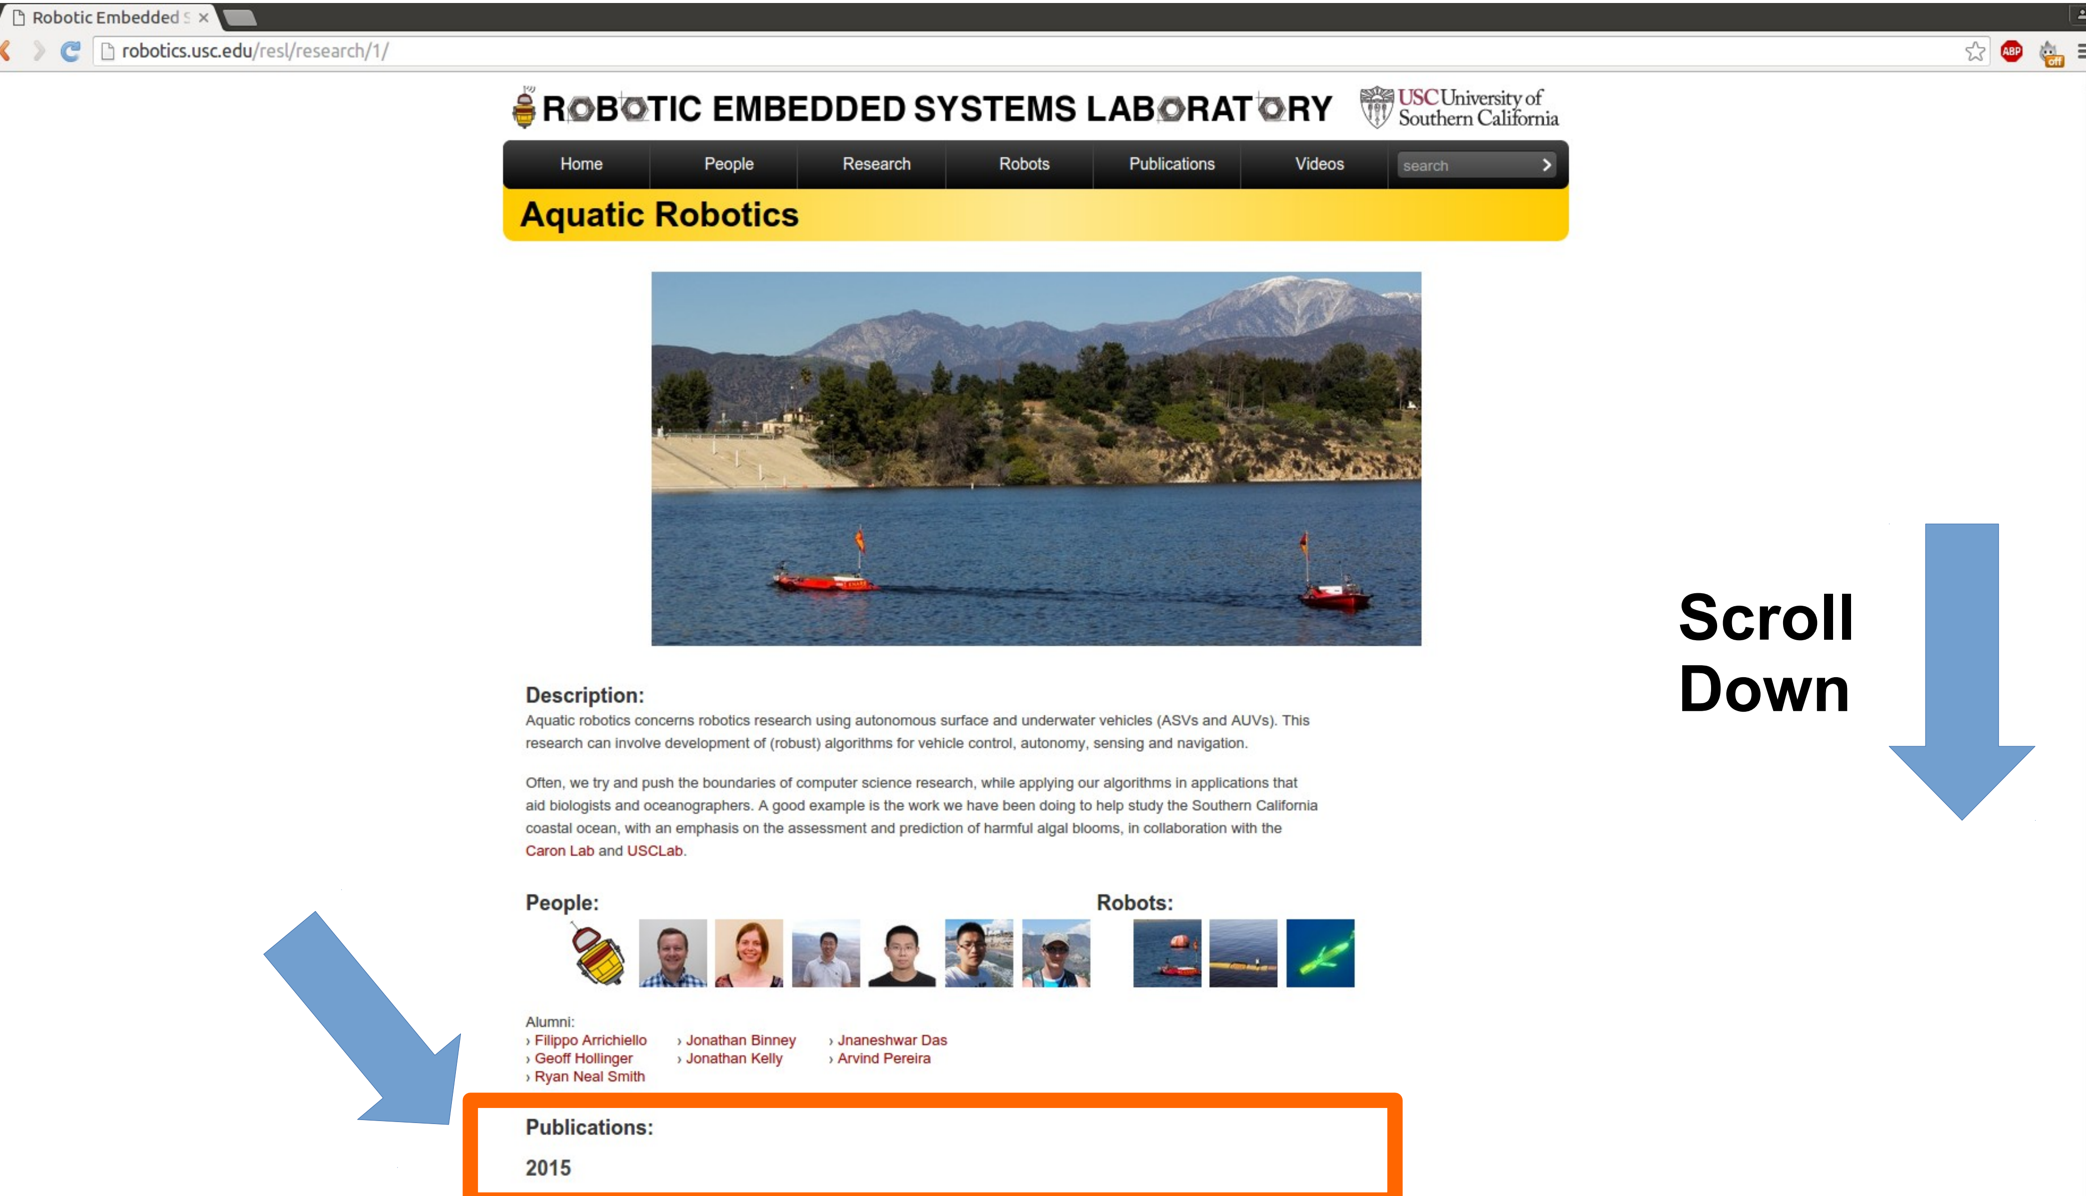

Scroll Down